

---

# The HP C1533A OEM Product Manual

**Edition 1.1GR  
May 1993**

---

**Part Number: C1533-90900**

**Hewlett-Packard Ltd.**  
Computer Peripherals Bristol  
Filton Road, Stoke Gifford  
Bristol BS12 6QZ, U.K.

# Notice

The information contained in this document is subject to change without notice.

**Hewlett-Packard makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this document.**

This document contains proprietary information which is protected by copyright. All rights reserved. No part of this document may be photocopied, reproduced or translated to another language without the prior written consent of Hewlett-Packard.

©Copyright 1993 by  
Hewlett-Packard Limited

## Printing History

---

Version	Date	Changed Pages
Edition 1	May 1993	All

---

# Contents

---

## Introduction

---

## Specifications Section

- Chapter S1* **Physical Specification**
  - Chapter S2* **Environmental Specification**  
containing requirements for installation and power
  - Chapter S3* **Reliability and Maintainability**  
containing reliability details and the service strategy and requirements
  - Chapter S4* **Performance Specification**  
containing details of capacity, transfer rate, timings, and retry limits
- 

## Operation and Installation Section

- Chapter O1* **Operation**  
containing all an operator needs to know to use and look after the drive, and how to interpret the front panel displays
  - Chapter O2* **Cartridges**  
describing DDS-logo media, the Media Recognition System, and how to take care of cartridges
  - Chapter O3* **Installation**  
containing instructions on how to install the drive, connecting to SCSI and power and setting the SCSI ID, and configuring data compression and the Media Recognition System
- 

## Interface Section

- Chapter I1* **Interface Implementation**  
listing supported SCSI messages and commands, describing typical SCSI operation, with details of disconnect and command queuing
- Chapter I2* **Messages**  
containing a summary of the messages supported, followed by details of the individual messages and status in numerical order
- Chapter I3* **Commands**  
containing an overview of the SCSI commands supported and pre-execution checks, followed by details of the individual commands in alphabetical order

---

## Theory of Operations Section

This section describes the mechanical and electrical operation of the drive down to a block level, and describes a typical signal path.

---

## Diagnostics Section

*Chapter D1* **Diagnostics Overview**  
explaining how to run diagnostic tests and obtain results

*Chapter D2* **Diagnostic Tests**  
giving details of the Power-Up and Self-Test diagnostic sequences

---

## Special Features Section

Special features are those extra facilities which set DDS-format drives apart from other tape drives.

*Chapter SF1* **Supporting DDS Features**  
explaining how to use DDS features, such as partitioning, fast-search, the Media Recognition System, and error-handling, and how to manage blocks and tape marks

*Chapter SF2* **Software Decompression**  
providing an example of host-controlled software decompression

---

## DDS Format Section

*Chapter DF1* **Digital Data Storage (DDS)**  
describing how the DDS-2 format builds on DAT, the layout of different tape areas, the concept of tape partitioning, and how the DDS-2 format extends DDS-1 in its use of entities

*Chapter DF2* **Fast-Search, Indexing, Error Correction**  
explaining how fast-search and indexing work in the drive, and the different methods of error correction which are available

---

## Data Compression Technique Section

*Chapter DC1* **Data Compression**  
describing the reasons for data compression, and the choice of the DCLZ algorithm

*Chapter DC2* **The DCLZ Algorithm**  
explaining how the DCLZ algorithm works in compressing and decompressing data

---

## Glossary

A more detailed list of contents is available at the start of each section.

---

---

# **Introduction**

---


# Introduction Contents

## The HP C1533A DDS-DC-Format Tape Drive

Product Features	1
The DDS-2 Format and Compatibility	2
The Purpose of this Manual	2
Related Documents	3

---

## The HP C1533A DDS-2 Tape Drive



The HP C1533A is a DDS-2 tape drive with a SCSI interface, designed for the OEM market. It conforms to the DDS-2 format standard for storing computer data, and incorporates a data compression capability using the DCLZ compression algorithm. It is a low-cost, high-capacity, high transfer-rate drive for storing data offline on tape. The drive fits a 3½-inch form factor.

The use of the DDS-2 format and data compression allow a 4–8 fold increase in capacity and transfer rate over that of a standard uncompressed drive using the DDS-1 format.

As a streaming tape drive, the HP C1533A implements all the mandatory and most of the optional SCSI commands for Sequential Access Devices.

---

## Product Features

The HP C1533A has the following features:

- A standard 3½-inch form factor
- A typical capacity of up to 8 gigabytes for a 120 m tape in a technical computing environment, and up to 16 gigabytes for a 120 m tape in a commercial environment
- A typical sustained streaming transfer rate of 510–1500 kilobytes/second
- A large 1 megabyte data buffer to maintain host transfer rate
- On-board DCLZ data compression and decompression subsystem
- Automatic error detection and correction
- Read-after-write error detection and correction
- Three levels of Error Correction Code (ECC)
- Synchronous or asynchronous communication modes
- Single-ended SCSI connection
- Electrically upgradeable firmware
- Compatibility with the SCSI-2 Sequential Access command set
- Support for disconnect and reselect
- Low cost

---

## The DDS-2 Format and Compatibility

The HP C1533A implements the DDS-2 (Digital Data Storage 2) format, an extension of the DDS-1 (originally called simply DDS) and DDS-DC formats, including data compression. DDS formats have been developed by Hewlett-Packard for data storage drives which use and build on DAT technology.

- The DDS-2 format allows compressed data to be stored in a way that maintains the full functionality of the DDS-1 format, and ensures backward compatibility with existing DDS drives.
- The HP C1533A can read any DDS tapes (DDS-1, DDS-DC or DDS-2), whether data compression is switched on or off.
- The drive automatically decompresses DCLZ data even if data compression is switched off.
- The drive supports the mixing of compressed and uncompressed data. This means that a tape can be labeled with uncompressed data and read in any DDS tape drive.

The DDS-2 format uses the same error correction techniques as the DDS-1 format. These incorporate the error correction techniques used in DAT drives with additional techniques specifically designed to provide the integrity of data necessary for computer use. The extra methods of error detection and correction are as follows:

- C3 ECC, a third level of Error Correction Code, providing correction across frames within groups.
- Read-After-Write, where data is verified immediately after it is written and re-written if there is any error.
- N-Group Writing, where every group is repeated a set number of times.
- Data Randomizer, for recoding the data in order to provide a consistent signal level. This enables read-after-write to be effective.
- Track checksums, for detecting and locating helical errors, and so improving the performance of C3 ECC.

Each tape can be formatted as a single data space, or as two partitions which can be written to independently.

DDS formats are used by numerous other tape drive manufacturers. Please consult your HP Sales Representative for the latest information about companies planning to build compatible drives.

For full details of the DDS-2 format, refer to the ECMA standard.

---

## The Purpose of this Manual

Together with the documents mentioned on the next page, this manual provides all the information about the HP C1533A which OEM manufacturers need in order to use the drive in their products.



---

## Related Documents

The following documents provide additional relevant information:

- **3,81 mm Wide Magnetic Tape Cartridge for Information Interchange—Helical Scan Recording—DDS-2 Format Using 120m Length Tapes, ECMA Standard**  
Available from the DDS Manufacturers Group or the address below
- **3,81 mm Wide Magnetic Tape Cartridge for Information Interchange—Helical Scan Recording—DDS Format, Standard ECMA 139**  
Available from the DDS Manufacturers Group or the address below
- **Designing a Data Storage Format for Digital Audio Tape (DAT)**  
Available from the DDS Manufacturers Group or the address below
- **Success with DDS Media**  
Available from the address below
- **Small Computer System Interface (SCSI-1), ANSI X3.131-1986**  
The ANSI authorized standard for SCSI implementation, available through ANSI
- **Enhanced Small Computer Systems Interface (SCSI-2)**  
ANSI X3T9.2-85-82 Rev.10, available through ANSI
- **Proposed American National Standard Data Compression Algorithm, Adaptive Coding with Embedded Dictionary, for Information Interchange**  
1st draft, 16 November 1990, X3B5/90-322, available through CBEMA, Washington DC
- **Proposed American National Standard Helical-Scan Digital Computer Tape Cartridge 3.81 mm (0.150 in.), Extended Digital Data Storage Format (DDS-DC), for Information Interchange**  
1st draft, 23 November 1990, X3B5/90-323, available through CBEMA

The first three of these documents are available from your HP Sales Representative, and also from:

**Veronica Neale**  
Hewlett-Packard Ltd.  
Filton Road  
Stoke Gifford  
Bristol BS12 6QZ  
UK

*Telephone:* +44 272 799910  
*Fax:* +44 272 236194  
*Telex:* 449206 HPCPB G



---

---

**Specifications**

---

# Specifications Contents

<b>S1</b>	<b>Physical Specification</b>	<b>1</b>
	<b>Dimensions</b>	<b>1</b>
	<b>Weight</b>	<b>2</b>
	<b>Clearance for Service</b>	<b>2</b>
	<b>Orientation</b>	<b>2</b>
	<b>Safety</b>	<b>2</b>
	<b>Conditions of Acceptability</b>	<b>2</b>
<b>S2</b>	<b>Environmental Specification</b>	<b>3</b>
	<b>Power Requirements</b>	<b>4</b>
<b>S3</b>	<b>Reliability and Maintainability</b>	<b>5</b>
	<b>Reliability</b>	<b>5</b>
	<b>Definition of Failure</b>	<b>5</b>
	<b>Mean Time Between Failures</b>	<b>5</b>
	<b>Product Life</b>	<b>5</b>
	<b>Service</b>	<b>6</b>
	<b>Regular Maintenance</b>	<b>6</b>
	<b>Field Replaceable Units</b>	<b>6</b>
	<b>Packaging for Return to the Factory</b>	<b>6</b>
<b>S4</b>	<b>Performance Specification</b>	<b>7</b>
	<b>Data Capacity</b>	<b>7</b>
	<b>Data Transfer Rate</b>	<b>7</b>
	<b>Sustained Data Transfer Rate to and from Tape</b>	<b>7</b>
	<b>Burst Data Transfer Rate to and from the SCSI Bus</b>	<b>7</b>
	<b>Data Interchange</b>	<b>7</b>
	<b>Load and Unload Times</b>	<b>8</b>
	<b>Error Rate</b>	<b>8</b>
	<b>Retries Limits on Rewrites</b>	<b>8</b>
	<b>Power-Fail Handling</b>	<b>8</b>

# Physical Specification


# S1

This section gives physical details of the HP C1533A. For details of cabling requirements for installation, see chapter O3 "Installation".

## Dimensions

Figure S-1 shows the general appearance of the drive.


Figure S-1



The HP C1533A  
DDS-Format Tape Drive

Figure S-2 gives details of the dimensions, tolerances and fixing centers for the drive.

Figure S-2



Drive Measurements

## Weight

One kilogram, without a cartridge.

## Clearance for Service


In the event of a fault, the entire drive is replaced. Consequently, it is recommended that you allow for easy access to the mounting screws (see figure S-2).

## Orientation

The HP C1533A can operate in three orientations, as shown in figure S-3. No adjustments are needed to internal settings for the different orientations.

Figure S-3

Operating orientations



## Safety

The HP C1533A conforms to the following safety standards:

- UL 1950—"Information Processing and Business Equipment" (Fifth Edition).
- CSA C22.2 Number 950, M1986—"Information Processing and Business Equipment".
- IEC 950—"Safety of Information Technology Equipment including Electrical Business Equipment" (First Edition).

## Conditions of Acceptability

The HP C1533A is for use only in equipment where the suitability of the combination has been determined by an appropriate certification organization (for example, Underwriters Laboratories Inc. or the Canadian Standards Association in North America, and the British Standards Institution or Verband Deutscher Elektrotechniker eV in Europe).

Other considerations include the following:

- 1 The drive must be installed in an enclosure to limit an operator's access to live parts, to provide system stability, and to give the drive the necessary grounding integrity.
- 2 The necessary voltage supplies must be provided. These supplies are Extra Low Voltage (SELV) for UL and CSA, or Safety Extra Low Voltage for BSI, VDE, and so on, of +5V and +12V dc.

---

**Note** The drive is only fused to protect it from excessive currents.

---

## 2 Specifications

# Environmental Specification

S2

These specifications are for the HP C1533A, not for any particular DDS-format tape. If a tape is used which has more stringent requirements than these specifications, the tape's requirements should be observed.

Parameter		Specification
Ambient Temperature	<i>Operating</i>	5°C to 40°C ( $\Delta T < 10^\circ\text{C/h}$ )
	<i>Non-Operating (mech)</i>	-40°C to 70°C ( $\Delta T < 20^\circ\text{C/h}$ )
	<i>Non-Operating (tape)</i>	-40°C to 45°C ( $\Delta T < 20^\circ\text{C/h}$ )
Humidity	<i>Operating</i>	20 to 80% RH, non-condensing Maximum wet bulb temperature = 26°C
	<i>Non-Operating (mech)</i>	5 to 95% RH ( $\Delta\text{RH} < 30\%/h$ )
	<i>Non-Operating (tape)</i>	20 to 80% RH ( $\Delta\text{RH} < 30\%/h$ )
Vibration	<i>Operating (3 Axes)</i>	<i>Swept Sine:</i> 0.3 g peak, 5–500 Hz @ 1 octave/min <i>Random:</i> 5–350 Hz @ 0.000533 g <sup>2</sup> /Hz 350–500 Hz @ -6 dB/octave 500 Hz @ 0.000271 g <sup>2</sup> /Hz (=0.5 g rms)
	<i>Non-Operating (3 Axes)</i>	<i>Swept Sine:</i> 0.75 g peak, 5–500 Hz @ 1 octave/min <i>Random:</i> 5–100 Hz @ 0.02 g <sup>2</sup> /Hz 100–137 Hz @ -6 dB/octave 137–350 Hz @ 0.0107 g <sup>2</sup> /Hz 350–500 Hz @ -6 dB/octave 500 Hz @ 0.0052 g <sup>2</sup> /Hz (=2.41 g rms)
Shock	<i>Operating (3 Axes)</i> <i>no performance change</i>	5.0 g peak for 3 ms — $v_2$ sine
	<i>Operating (3 Axes)</i> <i>no data loss</i>	8.0 g peak for 11 ms — $v_2$ sine
	<i>Non-Operating (3 Axes)</i> <i>no damage</i>	90.0 g peak for 3 ms — $v_2$ sine 30.0 g peak for 28 ms — trapezoidal
Altitude	<i>Operating</i>	0 km to 4.6 km (0 to 15,000 ft)
	<i>Non-Operating</i>	0 km to 15.2 km (0 to 50,000 ft)
Transportation	<i>Vibration</i> <i>Type 1 package</i> <i>(3 Axes)</i>	<i>Swept Sine:</i> 0.5 g peak @ 5–200–5 Hz @ 1 octave/min 5 min dwell at peak resonance <i>Random:</i> 5–100 Hz @ 0.015 g <sup>2</sup> /Hz 100–200 Hz @ -6 dB/octave 200 Hz @ 0.0038 g <sup>2</sup> /Hz (=1.47 g rms) (30 min/axis)
	Type 1: One drive packaged individually, also ten drives packaged together	<i>Impact</i> <i>Type 1 package</i> <i>Repetitive Shock</i> <i>Type 1 package</i>

Parameter		Specification
<b>Magnetic Field Susceptibility</b>	<i>Operating (no performance change)</i>	> 1 gauss @ 47.5–198 Hz (inside Helmholtz coil)
<b>ESD</b>	<i>Operating (no performance change)</i>	5–15 kV in 2.5 kV steps on the front panel only (ECC and re-read allowed)
	<i>Operating (no permanent damage)</i>	25 kV on the front panel only; includes IEC801-2 (EN55101-2) (power can be cycled, etc.)
<b>Radiated Susceptibility</b>	<i>Operating (no performance change)</i>	10 V/m over the frequency range 14 kHz–1 GHz
<b>Non-Operating Magnetic Field Emissions</b>	<i>Residual magnetic field emitted when packaged for transportation</i>	<0.525 microtesla @ 3 ft from any exterior surface of package (IATA transportation restrictions)
<b>Operating Magnetic Field Emissions</b>		<0.5 millitesla ac and/or dc on all exterior surfaces when operating in "worst-case" mode
<b>Radiated Emissions</b>		<FCC Level B when tested standalone
<b>Acoustic Noise Emissions</b>		<4.5 bel soundpower when tested standalone
<b>Thermal Profile</b>		Component surface temperatures not to exceed supplier's design limits under worst-case operating conditions. Nor shall any component be so hot as to be a potential hazard by burns or secondary reaction to fire.
<b>Basic Safety</b>		CSA Certified to C22.2 No. 950 UL Recognized to UL 1950 Complies with IEC 950 TUV mark to EN60950
<b>Suspended Particles</b>	<i>Operating and Non-Operating</i>	Suspended particle environment, particles <200 microgram/meter <sup>3</sup>

## Power Requirements

- +5V dc  $\pm$  5%, 1.2A typical
- +12V dc  $\pm$  10%, 0.2A typical (reprogramming EEPROMS requires 12V dc  $\pm$  5%)
- 0.6 amps max (<1 second)
- Power Consumption: <8.5W typical

The HP C1533A has internal fuses to protect it from a power supply delivering an excessive current to the drive. It is recommended that you route the two ground connections separately and directly through low impedance lines to the power supply unit, and connect them to a common grounding point there.



This chapter contains information regarding product reliability and service requirements.

---

## Reliability

### Definition of Failure

A failure is defined as any permanent malfunction of the drive that prevents you from retrieving data from a tape. This includes failure to power up, failure to unload or eject a cartridge, and failure to write and read data to and from the tape, and it assumes that both the drive and tape are being used within specification.

Faults are not considered failures when they are related to operator error, mishandling and abuse, system-related faults (cabling problems, unsupported systems, operating software, and so on), no trouble found, and transportation damage.

### Mean Time Between Failures

The mean time between failures (MTBF) for the HP C1533A is 200,000 power-on hours, where  $MTBF = \frac{\text{power on hours}}{\text{number of failures}}$  in the period.

Failures are only counted if the following conditions apply:

- Only DDS-approved cartridges are being used.
- The cartridges are neither damaged nor faulty.
- Environmental conditions (including DC voltage supplies) are within the specified limits.

---

**Note** Hewlett-Packard Company does not warrant that this predicted MTBF is representative of any particular unit installed for customer use. Failure rates are derived from a large database of test samples. The actual failure rate will vary from unit to unit.

---

### Product Life

The specified life of the drive is 10 years average, with a distribution of 8–12 years.

The usage model for this specification is as follows:

- 2 hours tape pulling time per day
- 6 days usage per week

This model is equivalent to 6,000 hours tape-pulling life, with a distribution of 4,800–7,200 hours.

---

**Note** The drive is designed to operate in streaming and repositioning modes. Sustained low data rates may cause premature failure due to an excessive number of starts and stops.

---

## Service

### Regular Maintenance

No regular maintenance is required apart from cleaning the heads after each 25 hours of use, or when the Caution Signal shows (see "Caution Signal" in chapter O1). Head cleaning is described in chapter O1.

### Field Replaceable Units

The drive is itself a single field replaceable unit. In the event of a problem with any of the sub-assemblies in the drive, the whole drive should be replaced, and the faulty drive returned to the factory.

### Packaging for Return to the Factory

The drive should be returned in its original packaging.

The Data Capacity, Data Transfer Rate and Error Rate specifications given in this chapter are based on the following assumptions:

- Media must conform to the DDS Media specification and not be damaged or faulty.
- The drive and media must remain within and always have remained within their respective operating and non-operating environmental specifications.
- The drive is at all times in good operating condition.
- The C3 ECC frame is generated on writing and used as necessary on reading.
- Read-after-write rewrites are used as necessary on writing.

## Data Capacity

The native (uncompressed) capacities of different length cartridges are as follows. Note that 60m and 90m cartridges are automatically written in DDS-1 format, and only 120m cartridges are written in DDS-2 format.

Length	Format	Native Capacity
60m	DDS-1	1.3 Gb typical
90m	DDS-1	2.0 Gb typical
120m	DDS-2	4.0 Gb typical

With data compression, capacity can be expected to increase in the ratio 2:1 for technical data, and 4:1 for standard data, though these figures depend on the exact nature of the data.

## Data Transfer Rate

### Sustained Data Transfer Rate To and From Tape

Sustained transfer rates are as follows:

- Native (uncompressed) mode: 510 Kb/s
- Typical compressed: 1.0 Mb/s
- Maximum compressed: 1.5 Mb/s

### Burst Data Transfer Rate To and From the SCSI Bus

The HP C1533A will sink and source data bursts to and from the SCSI bus at a maximum burst rate of 3 Mb/s using SCSI asynchronous transfers. Using synchronous transfers, the maximum burst rate is 7.5 Mb/s.

## Data Interchange

The drive is specified to interchange tapes with other drives which conform to the DDS-2 format. It can read tapes written by drives conforming to the DDS-1 and DDS-DC formats.

## Load and Unload Times

The load procedure takes 19 seconds from the time the cartridge is inserted into the HP C1533A to the time the drive is ready to start acting upon the next command from the host. Part of this time is used to access the System log at the start of the tape.

The time taken to unload and eject a tape depends heavily on the length of tape, the number of partitions and the current position on the tape. The tape must first be rewound, which could take up to 80s for a 120m tape if the current EOD is near the end of the tape. The drive then updates tape logs, and performs the unloading procedure. The result of this is that the entire eject sequence can take between 30s and 2 minutes.

## Error Rate

The unrecoverable error rate on read is specified as less than 1 in  $10^{15}$  bits read.

For further details, refer to "Designing a Data Storage Format for Digital Audio Tape (DAT)", produced for the DDS Manufacturers Group, and available from Hewlett-Packard (see Introduction).

## Retries Limits on Rewrites

For Read-After-Write error correction, each frame can be rewritten up to a maximum of 127 times (giving 128 writes of the frame). In addition, the total number of frame rewrites within one group is 128.

## Power-Fail Handling

If there is a power-fail, the HP C1533A performs the following actions, and reverts to its default configuration when power is restored:

- 1 The drive executes the Power-Up sequence of self-tests.
- 2 If a tape is present, it is threaded with the usual load sequence.
- 3 The drive fails the next host command and returns a CHECK CONDITION status. The host should then send a REQUEST SENSE command. The HP C1533A responds with the UNIT ATTENTION sense key, and sets the Additional Sense bytes to indicate that the drive has been power-cycled.

---

---

**Operation and Installation**

---

# Operation and Installation

<b>01</b>	<b>Operation</b>	
	<b>Status Display</b>	1
	<b>Caution Signal</b>	1
	<b>Diagnostic Display</b>	1
	<b>Operator Action</b>	2
	<b>Loading and Unloading Definitions</b>	2
	<b>Inserting Cartridges</b>	2
	<b>Removing Cartridges</b>	2
	<b>Forcing the Ejection of a Cartridge</b>	3
	<b>Write-Protecting a Cartridge</b>	3
	<b>Head Cleaning</b>	3
	<b>Power-Cycling</b>	4
	<b>Upgrading Firmware by Cartridge</b>	4
	<b>Avoiding Condensation</b>	4
<b>02</b>	<b>Cartridges</b>	
	<b>DDS-Logo Media</b>	5
	<b>The Media Recognition System</b>	5
	<b>Compatibility</b>	5
	<b>Ordering Tapes</b>	5
	<b>Write-Protecting a Cartridge</b>	6
	<b>Storage</b>	6
<b>03</b>	<b>Installation</b>	
	<b>Power Connection and Grounding</b>	7
	<b>SCSI Connection</b>	7
	<b>Setting the SCSI ID</b>	8
	<b>Termination Power</b>	8
	<b>Options Switches</b>	9

The front panel of the HP C1533A is used for the following:

- Inserting and removing cartridges
- Displaying drive and cartridge status

## Status Display

The front panel has two lights (LEDs) labeled "Tape" and "Clean". The Tape light shows green to show activity involving the tape, flashing slowly while the tape is being loaded or unloaded, steady when the tape is loaded, and fast flashing when read or write operations occur. The Clean light flashes amber slowly to indicate that a Caution Signal which either indicates that the heads need cleaning, or that the cartridge is near the end of its useful life. It shows a steady amber to indicate a hard fault.

The following diagram lists all the possibilities:

Figure O-1

	Tape	Clean	Meaning	Key
			Activity—load or unload	Off
			Activity—read or write	Green
			Cartridge loaded, drive online	Amber
			Caution signal	Flash Green ( $\frac{1}{2}$ s on, $\frac{1}{2}$ s off)
			Fault	Flash Amber ( $\frac{1}{2}$ s on, $\frac{1}{2}$ s off)
Front Panel Displays			Self-test in progress	Fast Flash Green ( $\frac{1}{4}$ s on, $\frac{1}{4}$ s off)

## Caution Signal

If an excessive number of read-after-write (RAW) or third level error correction (C3 ECC) errors are detected during normal operation, the Clean light flashes amber.

This condition could be caused by dirty heads, so the heads should be cleaned and the operation tried again. If the Caution Signal reappears, assume the cartridge is nearing the end of its useful life; copy the data on the cartridge onto a new one and discard the old cartridge. The Caution Signal is cleared after you have loaded a cleaning cartridge and the heads have been cleaned.

## Diagnostic Display

The purpose of the diagnostic firmware is to test the HP C1533A electronics for functionality. During power-up, the drive executes a self-test diagnostic sequence. This is shown by the Tape light flashing green at a rate of 1 cycle per second. If the self-test fails, the Clean light shows a steady amber.

---

## Operator Action

### Loading and Unloading Definitions

With regard to tapes, the drive may be in one of the following states:

- *Not Present* means there is no cartridge in the drive.
- *Present but not Loaded* means a cartridge is in the drive but not logically loaded. Logically loaded means that the drive has performed operations involving format evaluation (for example, reading the system log, read or write diagnostics), and is now prepared to accept commands which access the tape.
- *Loading* means the drive is loading the tape. This is a transition state between "Present but not Loaded" and "Loaded".
- *Loaded* means a cartridge is present in the drive and it has been logically loaded.
- *Unloading* means the drive is unloading the tape. This is a transition state between "Loaded" and "Present but not Loaded".
- *Ejecting* means the drive is in the process of ejecting the cartridge. This is a change of state from "Present but not Loaded" to "Not Present".

### Inserting Cartridges

Insert a cartridge into the slot on the front panel. As the cartridge is inserted, the drive takes it, performs a load sequence and goes online.

---

**Note** The drive can be configured to detect DDS Media Recognition System cartridges. If this is done and you load a cartridge which is not DDS Media Recognition System, the cartridge will be treated as write-protected; you will only be able to read it, not write to it. See Chapter O2 for more details.

---

The load sequence is as follows:

- 1 The drive mechanism threads the tape and rewinds to BOM (Beginning of Media). The tape is then reads the first Reference area on the tape (partition 1 on a 2-Partition tape) to find the tape format (DDS, audio, and so on) and the number of partitions.
- 2 If the tape is blank, the drive leaves the tape at BOP for Partition 0 and awaits the next command.
- 3 If the error rate is high the Caution Signal is set on the front panel lights.
- 4 The System area is then accessed and the Tape log read into the drive.
- 5 Finally the drive rewinds to BOP (partition 0) and goes online.

---

**Note** If the HP C1533A is power-cycled while a cartridge is loaded, the drive performs a load sequence and goes online.

---

After the load sequence, the drive sends a CHECK STATUS on receipt of the next SCSI command from the host. The UNIT ATTENTION key is set in the returned REQUEST SENSE data to indicate that the tape may have been changed.

### Removing Cartridges

A cartridge can be removed from the HP C1533A either in response to a SCSI LOAD/UNLOAD command with the Load bit set to 0, or by pressing the Unload button.

For either of these methods, if the host has not previously sent a PREVENT MEDIA REMOVAL command, the drive performs an Unload sequence. In this, the tape is rewound to BOP for Partition 0 and, if the tape is write-enabled, the copy of the Tape log held in



RAM is written back to tape. The tape is then rewound to BOM, unthreaded from the mechanism, and ejected.

---

**Note** If the host has previously sent a PREVENT MEDIA REMOVAL command, operation of the Unload button is disabled; pressing the button has no effect, and does not initiate an Unload sequence. The Unload button returns to normal operation following receipt of an ENABLE MEDIA REMOVAL command or a SCSI reset.

Also in this case, a LOAD/UNLOAD command with the Load bit set to zero will put the drive into the "present but not loaded" state, that is, it will take the drive offline and unthread the tape, but will not eject the cartridge from the drive.

---

## Forcing the Ejection of a Cartridge

---

**Caution** There is a possibility of losing data if you force an eject. The tape may also be invalidly formatted following the forced eject, because EOD may not have been written.

---

If you press the Unload button when the drive is busy, it may be a long time before the drive will respond to the request, because it will finish the task on which it is engaged first. This ensures that the task is terminated in a controlled manner and no data is lost.

However, there may be occasions when you need to unload a cartridge urgently, even at the risk of losing data. To do this, you must force an eject. There are two ways of doing this:

- Press the Unload button three times within a 5-minute period (or press the Unload button twice in the 5 minutes following a SCSI UNLOAD command).
- Hold the Unload button down for at least 5 seconds.

Following either of these actions, the drive waits until 35 seconds have passed from the time of the first press, to give the normal eject procedure a chance to be followed. Once this period of grace has elapsed, the tape is immediately unthreaded and the cartridge is ejected, regardless of what operation the drive was performing. The drive is then reset as though the power had been cycled.

Because the force eject may interrupt any operation, it is possible that the drive will not write EOD before the cartridge is ejected. This will cause loss of data, and result in a cartridge with an invalid format.

## Write-Protecting a Cartridge

Cartridges can be write-protected by sliding the tab on the rear of the cartridge so that the hole is open. In this state, data can be read from the tape but not written to it.

---

**Caution** The Tape log, which contains a history of usage of the tape, will not be updated when the cartridge is write-protected. It follows that the Tape log becomes inaccurate if a cartridge is used when write-protected.

---

## Head Cleaning

Users should clean the tape heads after every 25 hours of use. Cleaning the heads is also a sensible first step if the Caution Signal display (see earlier in this chapter) appears on the front panel lights.

In order to clean the heads, use a Cleaning Cartridge, HP 92283K. Insert the cleaning cartridge in the drive, when the drive will automatically load it and clean the heads. At the end, the cartridge is ejected. Note the date on the label on the cleaning cartridge, so that there is a record of how many times it has been used. After 25 uses, discard the cartridge.

## Power-Cycling

If the cartridge has not been ejected from the drive before power-cycling the unit, the tape will be "present but not loaded", and threaded when the drive is next powered up. A PREVENT MEDIA REMOVAL command sent before the power-cycle will be reset at power-up, so that the cartridge will now be ejected when an UNLOAD command is received.

If a bus reset occurs, or a Bus Device Reset is received, any buffered data will be written to tape before the drive goes into its reset sequence. This ensures data integrity.

---

## Upgrading Firmware by Cartridge

The HP C1533A allows you to upgrade its firmware by inserting a Firmware Upgrade tape. In order to do this, proceed as follows. The process will take at least 3 minutes to.

---

**Caution** Do not disconnect power to the drive at any time during the upgrade, or it may result in the drive having corrupt firmware.

---

- 1 Ensure that the drive is not engaged in any SCSI activity (the Tape light is lit steadily, not flashing).
  - 2 Insert the Firmware Upgrade tape into the drive.
  - 3 The upgrade process automatically takes place. About halfway through, the tape will be ejected. Remove it, and replace it in its box.
  - 4 You can be sure that the upgrade is complete when no activity can be seen on the front panel lights for at least 30 seconds.
- 

## Avoiding Condensation

Condensation is bad for tape drives where a thin tape is passing across a rapidly spinning head. To minimize the chance of condensation, adhere to the environmental requirements in the specifications and the guidelines in chapter O2.

To avoid temperature problems, observe these guidelines:

- Position the drive in a position where the temperature is relatively stable, for example, away from open windows, fan heaters, and doors.
- Avoid leaving cartridges in severe temperature conditions, for example in a car standing in bright sunlight.
- Avoid transferring data (reading from and writing to cartridges) when the temperature is changing by more than 10°C per hour.

## DDS-Logo Media

The need for reliability in storing computer data is greater than in recording music for reproduction. For this reason, the DDS logo was designed, indicating that a cartridge contains tape of data-grade quality. This logo is shown in figure O-2.

Although non-DDS cartridges may appear to work satisfactorily in the drive, their inferior specification threatens the integrity of your data. The Media Recognition System provides the drive with a means of detecting DDS cartridges.

Figure O-2



## The Media Recognition System

DDS Media Recognition System tapes have a series of stripes on the transparent leader tape at the beginning of the tape. The HP C1533A can detect these, and can be configured through the option switches on the underside to treat any tape without stripes as write-protected. In other words, it will only be allowed to read tapes which are not DDS Media Recognition System, not write to them.

You can recognize DDS Media Recognition System tapes by the logo in figure O-3.

Figure O-3

DDS Media  
Recognition  
System Logo

Tapes with the old DDS logo (figure O-2) are in no way inferior; they simply do not carry the stripes on the leader tape, so the drive cannot recognize them as DDS quality. All DDS cartridges produced from early 1993 should have the Media Recognition System stripes. HP cartridges have the same part numbers as before.

In order to configure the drive to respond to Media Recognition System tapes, see the section "Option Switches" in Chapter O3.

## Compatibility

The tape in DDS Media Recognition System has the same specification as tape in earlier DDS cartridges. Hewlett-Packard has used the opportunity to make small design changes to the cartridge shell, in order to make it resistant to even more severe shock and vibration when handled, but the cartridge and tape are wholly compatible with drives without the Media Recognition System.

Data on earlier DDS cartridges can still be read on the HP C1533A drive, whether the Media Recognition System is switched on or off. If you need to write to these tapes, turn the recognition system off as described in Chapter O3. The drive will then treat the tapes as it does DDS Media Recognition System tapes.

## Ordering Tapes

The HP C1533A can use any DDS Media Recognition System cartridges. However, for assured reliability, you are recommended to use Hewlett-Packard badged cartridges.

---

## Write-Protecting a Cartridge

Cartridges can be write-protected by sliding the tab on the rear of the cartridge so that the hole is open. In this state, data can be read from the tape but not written to it.

---

**Caution** The Tape log, which contains a history of usage of the tape, will not be updated when the cartridge is write-protected. It follows that the Tape log becomes inaccurate if a cartridge is used when write-protected.

---

---

## Storage


To minimize the chance of condensation, adhere to the environmental requirements in the specifications and the following guidelines:

- Only use cartridges at temperatures between 5°C (40°F) and 40°C (113°F). You can, however, store them at temperatures down to -40°C (-40°F).
- If you expose cartridges to temperatures outside the operating limits, stabilize them before you use them. To do this, leave the cartridges in the operating temperature for a minimum of two hours.

## Power Connection and Grounding


The power connections are shown in figure O-4.

Figure O-4



Grounding requirements are shown in figure O-5.

Figure O-5



## SCSI Connection

A single-ended SCSI configuration is supported. The hardware and termination specifications of this interface type can be found in Section 4, Physical Characteristics, of the SCSI X3T9/89-042 Standard.

Only unshielded connectors can be used.


The following are possible cable and terminator sources:

<b>Cable:</b>	28AWG Ribbon, Spectra Strip	132-2899-887 style 2697
<b>Connector:</b>	3M Connector	3425-7650
	3M Strain Relief	3448-3050

## Setting the SCSI ID

The SCSI ID is set using jumpers on the set of pins beside the SCSI connector at the rear of the drive, as shown in figure O-6.

Figure O-6



There are three significant bits in the ID, giving an ID range of 0 to 7 inclusive, as shown in the following table:

SCSI ID	BR 2	BR 1	BR 0	SCSI ID	BR 2	BR 1	BR 0	
0	0	0	0	4	1	0	0	1 = shorted
1	0	0	1	5	1	0	1	0 = open
2	0	1	0	6	1	1	0	
3	0	1	1	7	1	1	1	

The HP C1533A reads the SCSI ID at power-up and during self-test, in order to determine the selected target ID of the tape drive on the interface bus.


## Termination Power

The HP C1533A will provide termination power for the SCSI termination resistors if a jumper is set across the 2 pins marked "Term power" in figure O-6.

## Option Switches

The HP C1533A has a set of option switches on the underside of the unit, as shown in figure O-7.

Figure O-7



### Option Switches

Switches 1 and 2 are used to configure the data compression operation mode. The following table shows the available options.

Switch 1	Switch 2	Meaning
On	On	Compression enabled at power-on, with host control
On	Off	Compression enabled at power-on, no host control
Off	On	Compression disabled at power-on, the host is allowed to control compression
Off	Off	Compression disabled at power-on, no host control

Note that with switch 1 on, data written to the tape will be compressed without the knowledge of the host.

Switch 3 is used to configure the drive to respond to DDS Media Recognition System tapes:

Switch 3	Meaning
On	The Media Recognition System is disabled. This is the default. All DDS tapes will be treated the same, whether they possess the Media Recognition Stripes or not.
Off	The Media Recognition System is active. Non-Media Recognition System tapes are treated as if they are write-protected.

For more information on the Media Recognition System, see Chapter O2.

Switches 4 to 8 are used to specify drive connectivity and functionality according to host or customer requirements. The default setting is all switches on. Contact your HP distributor or dealer for further details.

Installation  
Option Switches

Switch 3 thru 8 Functionality

Host	Functions						Switches						
	MRS	Switch buffered	Performance	Caution Hard-error	Default fixed mode	Parity off	No Data phase discnt	8	7	6	5	4	3
HP		X		X				1	1	1	1	1	1
HP		X						1	1	1	0	1	1
HP			X	X				1	1	0	1	1	1
HP	X	X		X				1	1	1	1	1	0
HP	X	X						1	1	1	0	1	0
HP	X		X	X				1	1	0	1	1	0
DEC		X		X				0	1	1	1	1	1
DEC			X	X				0	1	0	1	1	1
DEC			X	X			X	0	1	0	1	0	1
DEC	X	X		X				0	1	1	1	1	0
DEC	X		X	X				0	1	0	1	1	0
DEC	X		X	X			X	0	1	0	1	0	0
SUN			X	X				0	0	0	1	1	1
SUN			X	X			X	0	0	0	1	1	0
SUN			X	X		X		0	0	0	1	0	1
SUN			X	X		X	X	0	0	0	1	0	0
SUN			X	X	X		X	0	0	0	0	1	0
SUN			X	X	X	X	X	0	0	0	0	0	0
SUN		X		X				0	0	1	1	1	1
SUN		X		X	X			0	0	1	0	1	1
SUN	X		X	X				1	0	0	1	1	1
SUN	X		X	X			X	1	0	0	1	1	0
SUN	X		X	X		X		1	0	0	1	0	1
SUN	X		X	X		X	X	1	0	0	1	0	0
SUN	X		X	X	X		X	1	0	0	0	1	0
SUN	X		X	X	X	X	X	1	0	0	0	0	0
SUN	X	X		X	X			1	0	1	0	1	1

Host connectivity options vary per the following Tables

Host	Mark Bit on Un-solicited sense	EW-EOM on read error	Write zero fm	No Valid Bit on 0 residue	Full reset	Signed residue
HP	1	1	1	1		
DEC					1	1
SUN		1				

Notes:

In the HP modes, the SCSI command following a LOAD command will result in a check condition, with the drive reporting via the Sense Key that the media may have been changed.

In the DEC/SUN modes, the check condition resulting from the LOAD operation is not issued.



### Special Function Definitions

**MRS:** When true, the Media Recognition System capabilities of the drive are enabled, allowing the drive to write data to MRS media only, and read from any 4mm media. When false, the drive is capable of reading and writing to any 4mm media.

**Mark Bit on Un-solicited Sense:** When true, the mark bit may be set on an un-solicited request sense (if tape located at a file/set mark). When false, this bit will never be set on an un-solicited request sense.

**EW-EOM on Read Error:** When true, EOM bit is set for check conditions resulting from read commands in the EW-EOM portion of the tape. When false, EOM is only set if at PEOP/M.

**Full Reset:** When true, drive runs through power-on selftests when a SCSI reset occurs. When false, the drive performs a "soft-reset" on the occurrence of a SCSI reset.

**Signed Residue:** When true, residue is negative on a backspace. When false, residue is absolute, positive, regardless of direction.

**Performance:** "Immediate" and "Infinite Flush" functions are selected when true. They are not selected when false.

**Immediate:** If selected, immediate mode is carried out on all commands. "switched buffer" = False is also enforced. The immediate flags on commands are silently ignored. If not selected, immediate flags are used, and "switched buffer" is left as-is.

**Infinite Flush:** When selected, write delay is set to 0. Write delay timing in mode select is silently ignored.

**Caution on Hard Error:** When true, if a hard error occurs on a read or write, then the LED display will display a caution signal (flashing yellow) indicating that the drive requires cleaning. When false, no special LED displays are used to report hard errors.

**Fixed mode:** When true, defaults to fixed mode with block size is set to 1024 bytes. When false, default block size is 0 (variable mode).

**No Parity:** When true, no parity checking of SCSI data is performed. When false, parity is checked.

**No Data Phase Disconnects:** When true, no disconnects in data phase. When false, disconnects are not regulated.

**Switch Buffered:** When set, or "Performance = True", defaults to buffered mode, and mode select can control buffering capability. When false, buffered mode is enforced and mode select buffer field is silently ignored.

**Write Zero FM:** When true, allows write zero filmarks on a write protected tape.

**No Valid Bit on 0 Residue:** When true, if check condition reported with the information bytes = 0 the valid bit will not be set. If false, the valid bit will be set even if the residue = 0.

Installation  
Option Switches

---

---

**Interface**

---

# Interface Contents

<b>I1</b>	<b>Interface Implementation</b>	
	<b>The HP C1533A SCSI Interface</b>	<b>1</b>
	Supported Messages	2
	Supported and Unsupported Commands	2
	<b>Signal Descriptions</b>	<b>3</b>
	<b>SCSI Bus Operation</b>	<b>4</b>
	Typical SCSI Operation	4
	Disconnect	5
	Control Byte—Flag and Link Bits	6
	Command Queuing	6
	Typical SCSI Sequences	7
<b>I2</b>	<b>Messages</b>	
	<b>Summary</b>	<b>9</b>
	<b>Message Details</b>	<b>9</b>
	COMMAND COMPLETE (00h)	
	EXTENDED MESSAGE (01h)	
	SAVE DATA POINTER (02h)	
	RESTORE POINTERS (03h)	
	DISCONNECT (04h)	
	INITIATOR DETECTED ERROR (05h)	
	ABORT (06h)	
	MESSAGE REJECT (07h)	
	NO-OP (NO OPERATION) (08h)	
	MESSAGE PARITY ERROR (09h)	
	LINKED COMMAND COMPLETE (0Ah)	
	LINKED COMMAND COMPLETE, with flag (0Bh)	
	BUS DEVICE RESET (0Ch)	
	IDENTIFY (80h–FFh)	
	<b>Status</b>	<b>13</b>
	GOOD (00h)	
	CHECK CONDITION (02h)	
	BUSY (08h)	
	INTERMEDIATE GOOD (10h)	
	RESERVATION CONFLICT (18h)	
<b>I3</b>	<b>Commands</b>	
	<b>Summary</b>	<b>15</b>
	<b>Command Details</b>	<b>15</b>
	<b>Pre-Execution Checks</b>	<b>16</b>
	Bad LUN Check	16
	Deferred Error Check	17
	Diagnostic Status Check	17
	Fixed Bit Check	17
	Flag Link Check	17
	Illegal Field Check	17
	Media Access Check	17
	Media Information Check	18
	Media Write Check	18

Parameter List Check	18
Reservation Check	19
Unit Attention Check	19
ERASE (19h)	20
INQUIRY (12h)	21
Inquiry Data Pages	23
LOAD/UNLOAD (1Bh)	25
LOCATE (2Bh)	27
LOG SELECT (4Ch)	29
LOG SENSE (4Dh)	31
Log Parameter Pages	33
Log Page Descriptor	33
Supported Log Pages Page	33
Writes and Read Error Counters Pages	34
Tape Log Page	36
Tape Capacity Log Page	37
Buffer Trace Log Page	39
Device Trace Log Page	40
Writes and Read Frame Error Counter Pages	41
Bad Group Log page	42
Drive Counters Log Page	43
Mechanism Counters Log Page	44
Data Compression Transfer Log Page	45
Data Compression Trace Log	46
Channel Trace Log	47
MODE SELECT (15h)	48
Mode Parameter Pages	50
Mode Page Representation	50
Mode Data Format	51
Mode Block Descriptor	52
Disconnect-Reconnect Page	52
Data Compression Characteristics Page	54
Device Configuration Page	56
Medium Partitions Parameter Page	58
MODE SENSE (1Ah)	60
PREVENT/ALLOW MEDIUM REMOVAL (1Eh)	62
READ (08h)	63
READ BLOCK LIMITS (05h)	65
READ BUFFER (3Ch)	66
READ POSITION (34h)	68
RECEIVE DIAGNOSTIC RESULTS (1Ch)	70
RELEASE UNIT (17h)	71
REQUEST BLOCK ADDRESS (02h)	72
REQUEST SENSE (03h)	73
Request Sense Data	73
Sense Data Management	76
Current Sense	76
Unit Attention Sense	77
Deferred Error Sense	78
Sense Keys and Additional Sense Codes	79

## Interface Contents

RESERVE UNIT (16h)	84
REWIND (01h)	85
SEEK BLOCK (0Ch)	86
SEND DIAGNOSTIC (1Dh)	87
SPACE (11h)	88
TEST UNIT READY (00h)	90
VERIFY (13h)	91
WRITE (0Ah)	93
WRITE BUFFER (3Bh)	95
WRITE FILEMARKS (10h)	97

The HP C1533A DDS-2 drive uses SCSI-2 as the interface to connect to the host system.

---

**Note** Throughout this section, the term SCSI implies SCSI-2.

---

This section covers the implementation of SCSI in the drive.

- Chapter I1, this chapter, gives an overview of SCSI in the HP C1533A, a list of supported and unsupported commands, and a description of SCSI operation with typical sequences.
- Chapter I2 provides details of all the SCSI messages and Status conditions that may be generated. This covers section 6 of the ANSI standard.
- Chapter I3 gives details of all the SCSI commands that are supported.

This section, together with the SCSI documents listed in the Introduction, provides the information required to integrate the HP C1533A with SCSI-compatible computer systems.

See also the chapter on "Special Features".

## The HP C1533A SCSI Interface

The Small Computer System Interface (SCSI) is an industry standard interface, approved by the American National Standards Institute (ANSI). It is recommended that you read the ANSI standard document in conjunction with this manual. The ANSI specification defines the interface in general while this document describes the HP C1533A implementation.

The HP C1533A SCSI implementation provides a drive with a standard set of features and functions. These include the following:

- Synchronous data transfers at burst speeds up to 7.5 MHz.
- Asynchronous data transfers at burst speeds up to 3 MHz
- Single-ended glitchless drivers
- Single-ended SCSI connection
- 8-bit wide data transfers plus parity
- Ability to retry SCSI command, message, status and data phases (such as response to parity errors or errors detected by the host)
- Multi-initiator support (queuing of requests, maintenance of separate sense information)
- Optional termination power
- Implementation of all mandatory and most optional commands of the Sequential Access command set

## Supported Messages

The following messages are supported by the HP C1533A:

- ABORT
- BUS DEVICE RESET
- COMMAND COMPLETE
- DISCONNECT
- EXTENDED MESSAGE—SYNCHRONOUS DATA TRANSFER REQUEST
- IDENTIFY
- INITIATOR DETECTED ERROR
- LINKED COMMAND COMPLETE
- LINKED COMMAND COMPLETE (with flag)
- MESSAGE PARITY ERROR
- MESSAGE REJECT
- NO-OP (NO OPERATION)
- RESTORE POINTERS
- SAVE DATA POINTER

For implementation details on these messages, see chapter I2, "Messages".

## Supported and Unsupported Commands

The following commands are supported by the HP C1533A. They include all Mandatory and Extended commands and most Optional commands.

- |                               |                         |
|-------------------------------|-------------------------|
| ■ ERASE                       | ■ RELEASE UNIT          |
| ■ INQUIRY                     | ■ REQUEST BLOCK ADDRESS |
| ■ LOAD/UNLOAD                 | ■ REQUEST SENSE         |
| ■ LOCATE                      | ■ RESERVE UNIT          |
| ■ LOG SELECT                  | ■ REWIND                |
| ■ LOG SENSE                   | ■ SEEK BLOCK            |
| ■ MODE SELECT                 | ■ SEND DIAGNOSTIC       |
| ■ MODE SENSE                  | ■ SPACE                 |
| ■ PREVENT/ALLOW MEDIA REMOVAL | ■ TEST UNIT READY       |
| ■ READ                        | ■ VERIFY                |
| ■ READ BLOCK LIMITS           | ■ WRITE                 |
| ■ READ BUFFER                 | ■ WRITE BUFFER          |
| ■ READ POSITION               | ■ WRITE FILEMARKS       |
| ■ RECEIVE DIAGNOSTIC RESULTS  |                         |

For implementation details on these commands, see chapter I3, "Commands".

The following Optional commands are not supported by the HP C1533A.

- |                     |                         |
|---------------------|-------------------------|
| ■ CHANGE DEFINITION | ■ COPY AND VERIFY       |
| ■ COMPARE           | ■ READ REVERSE          |
| ■ COPY              | ■ RECOVER BUFFERED DATA |



---

## Signal Descriptions

The HP C1533A SCSI interface consists of eighteen signals—nine control lines and nine data lines. A description of these signals is given in the following table.

Signal	Name	Driven by	Description
-BSY	Busy	—	OR-tied signal used to indicate that the SCSI bus is in use.
-SEL	Select	Initiator Target	Used to select a Target during the Selection phase. Used to select an Initiator during the Reselection phase.
-C/D	Control/Data	Target	Indicates whether Control or Data information is on the data bus. True (low) indicates Control, and false (high) indicates Data.
-I/O	Input/Output	Target	Controls the direction of data movement on the bus with respect to the Initiator. This signal is also used to distinguish between Selection and Reselection phases. True (low) indicates input to the Initiator, and false (high) indicates output from the Initiator.
-MSG	Message	Target	Indicates a Message phase on the bus. True (low) indicates a Message phase, and false (high) indicates Command, Data or Status phases.
-REQ	Request	Target	Indicates a request for a REQ/ACK data transfer handshake.
-ACK	Acknowledge	Initiator	Indicates an acknowledgement for a REQ/ACK data transfer handshake.
-ATN	Attention	Initiator	Indicates that the Initiator has a message to send to the Target.
-RST	Reset	—	OR-tied signal which is used to indicate a Reset condition.
DB(7-0)	Data Bus	—	8 data-bit signals, which with the parity-bit signal form the data bus. DB7 is the Most Significant Bit, and has the highest priority during the Arbitration phase.
DB(P)	Data Bus	—	Data parity bit which is set to <i>odd</i> , but is a jumper selectable option. Parity is not valid during the Arbitration phase.

---

## SCSI Bus Operation

### Typical SCSI Operation

This example describes the typical SCSI bus sequence between the host and the target.

- 1 The host arbitrates for the SCSI bus by asserting BSY and the data line corresponding to its bus ID. If any other devices wish to compete for the bus, they also assert BSY and the appropriate data line. Each arbitrating device then inspects the data bus and the device with the highest ID wins it. All the other devices must release BSY and their data lines.
- 2 The host attempts to select the target by asserting SEL and releasing BSY. The host maintains its ID and asserts the target's ID on the data bus. Each target then checks the data lines. If the target's ID matches that on the data bus, it accepts selection by asserting BSY. Once the host has detected BSY being asserted, it asserts ATN to indicate that it will want the target to go to the MESSAGE OUT phase. The host releases SEL.
- 3 The target now has control of the SCSI bus and it is the target which switches between phases. The target responds to the ATTENTION condition and initiates the MESSAGE OUT phase. The host sends an IDENTIFY message which tells the target which logical unit the host wishes to talk to. The fact that the target responds to the ATN indicates to the host that the target can accommodate more than just a COMMAND COMPLETE message.
- 4 The target initiates the COMMAND phase and transfers the Command Descriptor Block from the host. In the COMMAND phase, the target decodes the command and either executes the command (TEST UNIT READY) or sets itself up for a data transfer to the host (for example, READ, WRITE, INQUIRY). The target then either switches to the STATUS phase if the command is complete, or the DATA phase if it is ready to transfer data.
- 5 The data transfer length is set by the host in the Command Descriptor Block. The target remains in the DATA phase until all the data is transferred.
- 6 The target then initiates a STATUS phase and transfers one byte to the host to indicate whether it has completed the command successfully. If the target has detected an error, the next command that the host is expected to send is REQUEST SENSE. This allows the target to return further status information to the host.
- 7 The target completes the SCSI sequence by going to the MESSAGE IN phase and transferring a COMMAND COMPLETE message to the host. The target then releases BSY, allowing the bus to go to the BUS FREE state.

## Disconnect

In order to improve bus usage and performance, the HP C1533A is capable of disconnecting from the host in order to free the bus to allow other requests to be sent to other targets. To do so, however, the host needs to support Disconnect/Reselect. If Disconnect is implemented, the procedure is as follows:

- 1 The host arbitrates for the SCSI bus and if it wins it, selects the target device. Before releasing SEL and completing the selection phase, the host asserts the ATN line. The host then releases SEL and BSY. The target now has control of the SCSI bus. By asserting ATN, the host has indicated that the target should go to a MESSAGE OUT phase.
- 2 After the SELECTION phase is completed, the target responds to the host's ATTENTION condition by initiating a MESSAGE OUT phase. It receives a message from the host which tells it whether the host can support Disconnect/Reselect and the desired logical unit number on the target.
- 3 The I/O activity from this point is controlled entirely by the target. The target initiates the COMMAND phase and reads in the Command Descriptor Block from the host. After decoding the command, the target determines whether it should disconnect from the bus. The target disconnects from the bus for any non-trivial commands.
- 4 The disconnect process is when the target initiates a MESSAGE IN phase and sends the host a SAVE DATA POINTERS (during a DATA phase only) and a DISCONNECT message. Following the MESSAGE IN phase, the target releases BSY, freeing the bus which then enters the BUS FREE state. The host can now select another target, or allow another target to win the bus and reselect the host.
- 5 Although the host and the target are physically disconnected, they are still logically connected. Both know that they have a command to finish and will return to that job later. This principle allows many I/O commands to be executed simultaneously using a single peripheral bus. Once the target has completed a task and is ready to communicate with the host, it must re-establish the physical path. The reselection process involves the target arbitrating for the bus and reselecting the host. After the physical reconnection is made, the target sends an IDENTIFY message to the host to indicate which target logical unit it is. Following completion of this phase, the target initiates the next appropriate phase for the command, usually a DATA phase.
- 6 During a large data transfer, the target may disconnect at intervals depending on its use of the bus. The drive optimizes its use of the bus so as to maximize the transfer rate when it is connected to the host, and to minimize the time for which it holds the bus without handshakes. If the target disconnects, during a data transfer, the target initiates a MESSAGE IN phase and send the host a SAVE DATA POINTERS message and a DISCONNECT message. The host responds to the SAVE DATA POINTERS message by saving the current data pointer. After transmission of the DISCONNECT message the target releases BSY, freeing the bus.
- 7 Once the target is again ready to reselect the host, it goes through the same process as before—arbitrating for the bus, reselecting the host and sending an IDENTIFY message. However, the host's response is slightly different in this case since the disconnect was during a data transfer. Host acceptance of the IDENTIFY message also implies a RESTORE DATA POINTERS message to the host. The data transfer can now be resumed.
- 8 After completion of the data transfer, the target initiates a STATUS phase and sends a single status byte to the host. The final action of the target is to initiate a MESSAGE IN phase and send a COMMAND COMPLETE message to the host.

---

**Note** The HP C1533A will disconnect on completion of a data transfer if the final transfer occurs on a disconnect boundary, before initiating the STATUS phase. This is intended to optimize bus usage.

The HP C1533A also disconnects to queue commands from other initiators. See "Command Queueing" below.

---

The HP C1533A does not disconnect on receipt of the following commands:

- INQUIRY
- TEST UNIT READY
- REQUEST SENSE

The drive will disconnect on other commands if it is configured to do so.

### Control Byte— Flag and Link Bits

The control byte of the Command Descriptor Block contains the Flag and Link bits. Use of these bits is entirely host-dependent. Setting the Link bit provides an automatic link to the next command. Instead of returning the usual COMMAND COMPLETE message and going to BUS FREE, the drive goes straight to the COMMAND phase and handshakes in the next Command Descriptor Block from the host. If at any time an error is detected, the drive generates a CHECK CONDITION which breaks the sequence of Linked commands. The host is then expected to issue a REQUEST SENSE command to find the cause of the error.

The following table shows the different permutations of Flag and Link settings.

Flag	Link	Status	Message or Sense Data
0	0	—	—
0	1	INTERMEDIATE GOOD	LINKED COMMAND COMPLETE
1	0	CHECK CONDITION	ILLEGAL COMMAND DESCRIPTOR sense key
1	1	INTERMEDIATE GOOD	LINKED COMMAND COMPLETE with Flag

The host may use the Flag bit to trigger an internal interrupt or take some specific action following completion of the Flagged command within a Linked sequence. The use of the Link and Flag bits means that the drive holds the bus continuously. The host should not enable disconnect while using Linked commands because it would have to arbitrate and reselect the target between commands and so lose the performance gain Linked commands can give.

### Command Queueing

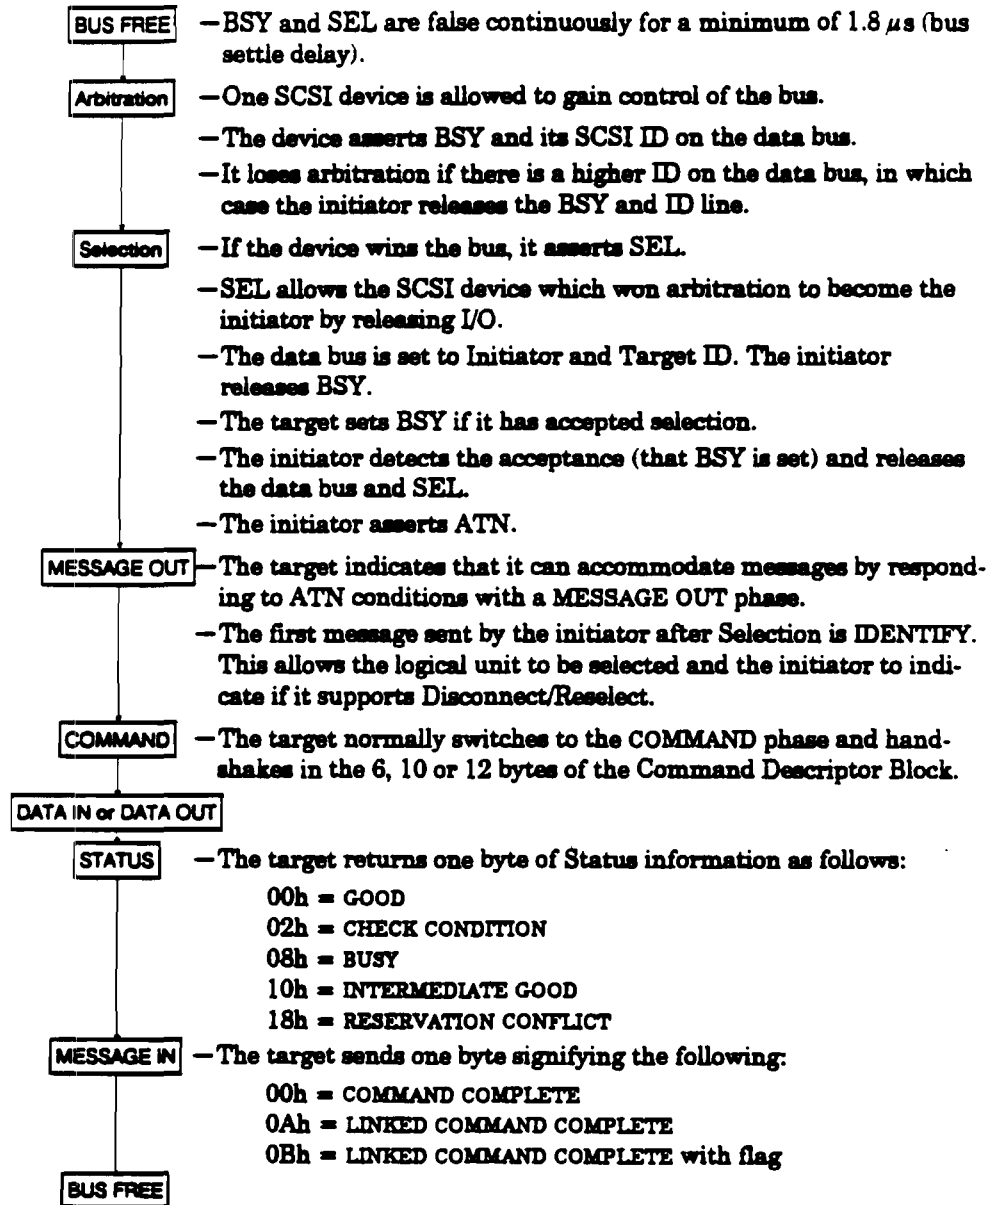
The HP C1533A does not support command queueing as defined in the SCSI-2 standard. For the drive, queueing works as follows:

If the drive is executing a command, it will disconnect and queue commands sent from other hosts, providing its current host will permit disconnection and reselection. Once a command is queued, the HP C1533A will reselect and continue executing the command at the head of the queue. Commands are queued in the order they arrive, and executed in the same order.

The HP C1533A will not queue a second command from a host for which it already has an active operation. The drive will return CHECK CONDITION status with the sense key set to ABORTED COMMAND, and additional sense code of 4E00h (overlapped commands attempted).

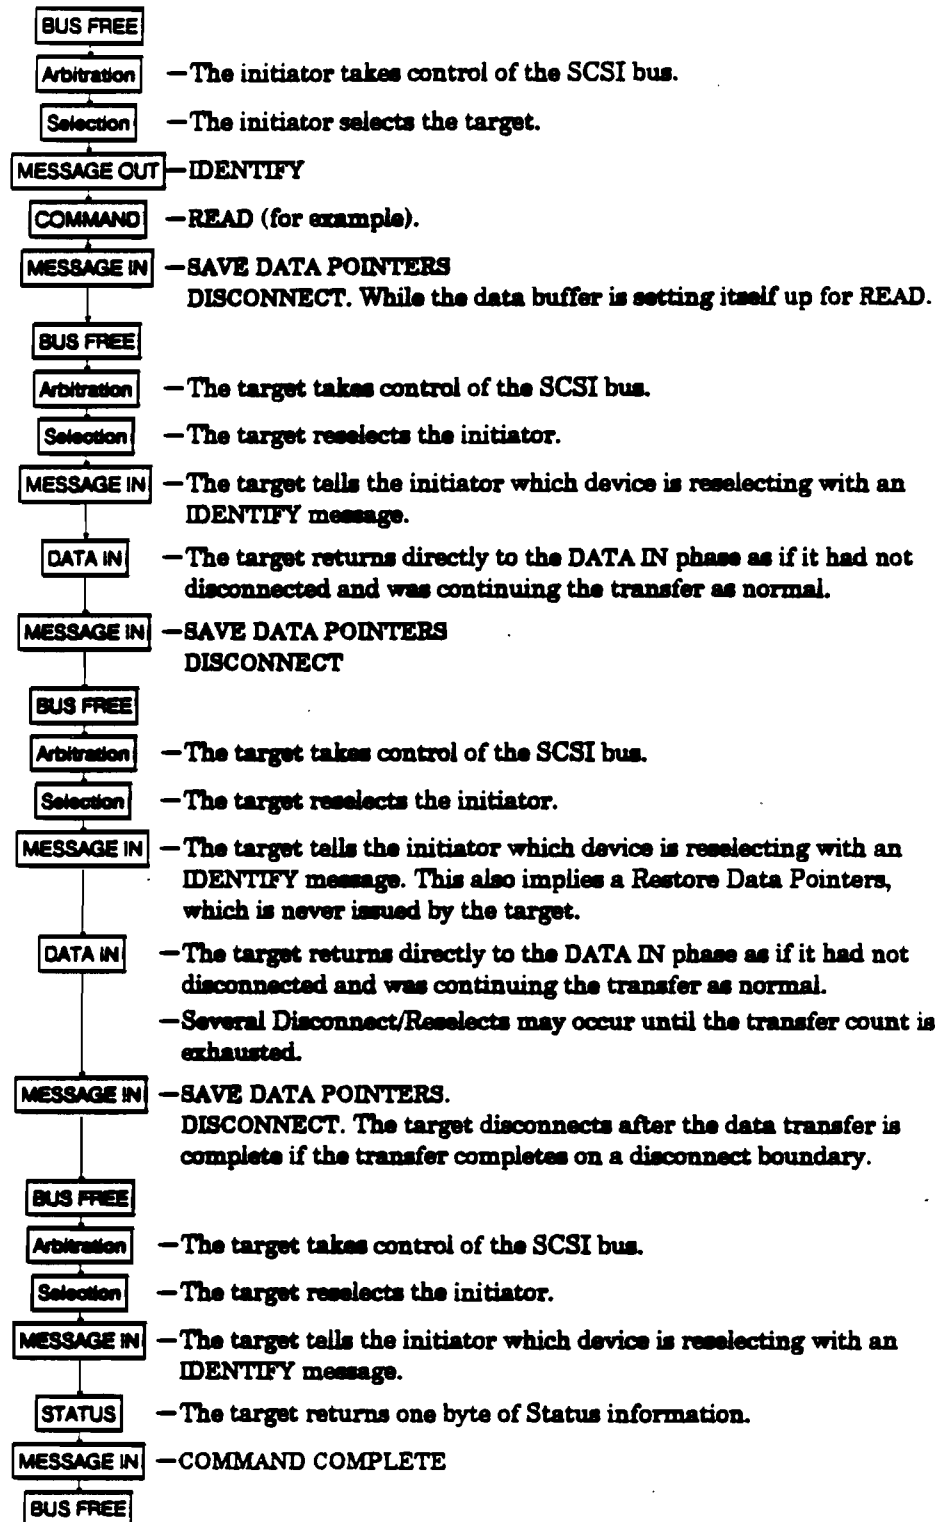
## Typical SCSI Sequences

### 1 Without Disconnect



## SCSI Bus Operation

## 2 With Disconnect:



This chapter includes all the HP C1533A SCSI messages, both supported and unsupported. Parts of this chapter come from Section 5, Logical Characteristics, of the SCSI X3T9/89/042 standard.

## Summary

The message system provides an initiator and a target on the SCSI bus with a means of managing communication. The available messages are as follows:

Code (hex)	Description	Supported
00	COMMAND COMPLETE	Yes
01	EXTENDED MESSAGE	Partially
02	SAVE DATA POINTER	Yes
03	RESTORE POINTERS	Yes
04	DISCONNECT	Yes
05	INITIATOR DETECTED ERROR	Yes
06	ABORT	Yes
07	MESSAGE REJECT	Yes
08	NO-OP	Yes
09	MESSAGE PARITY ERROR	Yes
0A	LINKED COMMAND COMPLETE	Yes
0B	LINKED COMMAND COMPLETE (with flag)	Yes
0C	BUS DEVICE RESET	Yes
0D-7F	RESERVED CODES	No
80-FF	IDENTIFY	Partially

## Message Details

### COMMAND COMPLETE (00h)

A target sends **COMMAND COMPLETE** to an initiator to indicate that the execution of a command (or a series of linked commands) is complete and that valid status has been sent to the initiator. After it has sent the message successfully, the target goes to the Bus Free phase by releasing **BSY**.

If the drive receives **COMMAND COMPLETE** as a target, it handles it as an illegal message. The drive returns **MESSAGE REJECT** and enters the Status phase reporting **CHECK CONDITION** with the sense key set to **COMMAND ABORTED**.

## EXTENDED MESSAGE (01h)

EXTENDED MESSAGE is usually sent as the first byte of a multiple-byte message by either initiator or target. It has the following format:

Byte 0	Extended Message Identifier—01h
Byte 1	Extended Message Length
Byte 2	Extended Message Code
Byte 3	Extended Message additional parameters

The Extended Message Length gives the length of the Extended Message plus the number of additional parameters that are to follow. The length does not include bytes 0 and 1.

The only Extended Message Code supported by the HP C1533A is 01h, Synchronous Data Transfer Request (SDTR). Any other message code is handled as an illegal message; the drive returns MESSAGE REJECT and enters the Status phase reporting CHECK CONDITION with the sense key set to COMMAND ABORTED.

An SDTR message has the following format:

Byte 0	01h
Byte 1	03h
Byte 2	Synchronous Data Transfer Request code—01h
Byte 3	Transfer Period x 4 nanoseconds
Byte 4	REQ/ACK offset

The HP C1533A supports both *synchronous* and *asynchronous* data transfer.

The HP C1533A will initiate an SDTR before or after a command has been read, when the current data transfer mode is undefined. The default is to negotiate before the Command phase. However, the drive will initiate negotiation *after* the Command phase in the following circumstances:

- The firmware has been configured not to negotiate before the Command phase.
- A host-initiated negotiation, which occurred before the Command phase of a previous operation, failed.

If the host supports synchronous data transfer, it will respond with an SDTR. If the host does not respond to the SDTR, or sends a MESSAGE REJECT, the HP C1533A will configure all future data transfers with this host to *asynchronous* mode. If the HP C1533A receives a delayed SDTR response from the host (548 nanoseconds), the drive will send a MESSAGE REJECT and assume that data transfers are to be in *asynchronous* mode.

If the host sends an SDTR when the drive has not, the drive will respond with a message-in-SDTR and subsequent data transfers will take place in *synchronous* mode.

The SDTR message from the host must be the first message after a selection phase, otherwise it is handled as an illegal message. In this case, the drive returns MESSAGE REJECT and enters the Status phase reporting CHECK CONDITION with the sense key set to COMMAND ABORTED.

## SAVE DATA POINTER (02h)

A target sends SAVE DATA POINTER to the initiator to save a copy of the present active Data pointer for the logical unit which is currently attached.

The HP C1533A accepts this message when it is the initiator. As a target, the drive sends this message before a disconnect during a data transfer. It does not send a SAVE DATA POINTER message if it intends to move directly to Status phase.

When received as a target, it is handled as an illegal message; the drive returns MESSAGE REJECT and enters the Status phase reporting CHECK CONDITION with the sense key set to COMMAND ABORTED.



## RESTORE POINTERS (03h)

A target sends RESTORE POINTERS to direct the initiator to restore the pointers for the currently attached logical unit which were saved most recently to the active state. Pointers to the command, data, and status locations for the logical unit are restored to the active pointers. Command and Status pointers are restored to the beginning of the present Command and Status areas. The Data pointer is 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 for that logical unit.

As a target, the HP C1533A does not normally send a RESTORE POINTERS message. Rather, it relies on the Restore Pointers implicit in a reselection. (RESTORE POINTERS may be sent when retrying a SCSI phase).

When acting as initiator, the HP C1533A accepts a RESTORE POINTERS (or a reselection without a RESTORE POINTERS) if SAVE DATA POINTER was received earlier.

When received as a target, RESTORE POINTERS is handled as an illegal message; the drive returns MESSAGE REJECT and enters the Status phase reporting CHECK CONDITION with the sense key set to COMMAND ABORTED.

## DISCONNECT (04h)

A target sends DISCONNECT to tell an initiator that the present physical path is going to be broken—the target intends to disconnect by releasing BSY. Later, reselection is required in order to complete the current operation.

---

**Note** The initiator detects a catastrophic error condition if the Bus Free phase occurs (other than as result of a reset condition) without first receiving a DISCONNECT or COMMAND COMPLETE message. If the target intentionally creates this condition, the target clears the current command. This message does not cause the initiator to save the Data pointer.

---

The HP C1533A supports this message as an initiator.

As a target, the HP C1533A supports this message and disconnects if the initiator allows it. The Disconnect strategy of the drive is discussed in Chapter 11.

When received as a target, DISCONNECT is handled as an illegal message; the drive returns MESSAGE REJECT and enters the Status phase reporting CHECK CONDITION with the sense key set to COMMAND ABORTED.

## INITIATOR DETECTED ERROR (05h)

An initiator sends INITIATOR DETECTED ERROR to inform a target that an error has occurred which does not prevent the target from trying the operation again—for example, a parity error. At this time, present pointer integrity is not assured. A RESTORE POINTERS message or a disconnect followed by a reselection, will restore the pointers to their previously defined state.

As initiator, the HP C1533A does not send this message. If the drive receives this message for any phase other than Data-in or Status, it enters the Status phase, reporting CHECK CONDITION with the sense key set to COMMAND ABORTED.

## ABORT (06h)

An initiator sends ABORT to a target to clear the present operation. All pending data and status from the affected logical unit is cleared, and the target goes to the Bus Free phase. This message can be sent to a logical unit that is not currently performing an operation for the initiator.

Messages  
Message Details

This message can also be sent to a logical unit which is currently performing an operation for the host. It should be sent using "select to abort", whereby the host selects and sends an IDENTIFY-ABORT sequence.

### MESSAGE REJECT (07h)

Either initiator or target sends MESSAGE REJECT to indicate that the last message received was inappropriate or has not been implemented.

In order to indicate that it is about to send this message, the initiator asserts the ATN signal before it releases ACK for the REQ/ACK handshake of the message that is to be rejected. MESSAGE REJECT is issued in response to any messages which the drive considers to be illegal or not supported. The illegal message causes the current command to be aborted, and the drive enters the Status phase reporting CHECK CONDITION with the sense key set to COMMAND ABORTED.

When the drive is acting as a target and receives the message, if the drive was previously in Message In phase (that is, sending messages to the host), MESSAGE REJECT causes the drive to go straight to the Status phase with a CHECK CONDITION status and a sense key of COMMAND ABORTED. If the drive was not previously in Message In phase, MESSAGE REJECT is handled as an illegal message.

### NO-OP (NO OPERATION) (08h)

If a target requests a message, the initiator sends NO-OP if it does not currently have any other valid message to send. The message is accepted when the drive is acting as a target, and may be sent when it is an initiator. If a NO-OP is received during any phase, the drive is likely to repeat that phase.

### MESSAGE PARITY ERROR (09h)

An initiator sends MESSAGE PARITY ERROR to a target to indicate that one or more bytes in the last message it received had a parity error.

To indicate that it intends to send the message, the initiator sends the ATN signal before it releases ACK for the REQ/ACK handshake of the message that has the parity error. This provides an interlock so that the target can determine which message has the parity error.

As initiator, the HP C1533A does not send this message. If the drive receives the message while acting as a target, it enters the Status phase reporting CHECK CONDITION with the sense key set to COMMAND ABORTED.

### LINKED COMMAND COMPLETE (0Ah)

A target sends LINKED COMMAND COMPLETE to an initiator to indicate that the execution of a linked command is complete and that status has been sent. The initiator then sets the pointers to the initial state for the next linked command.

The HP C1533A returns this message to the host when a linked command is completed and the Flag bit in the command descriptor block is not set.

When it receives the message as a target, the drive handles it as an illegal message; it returns MESSAGE REJECT and enters the Status phase reporting CHECK CONDITION with the sense key set to COMMAND ABORTED.

### LINKED COMMAND COMPLETE, with Flag (0Bh)

A target sends LINKED COMMAND COMPLETE with Flag to an initiator to indicate that the execution of a linked command with the Flag bit set to one is complete and that status has been sent. The initiator then sets the pointers to the initial state of the next linked

command. Typically, this message would be used to cause an interrupt in the initiator between two linked commands.

When the HP C1533A is acting as a target, it handles this as an illegal message; the drive returns MESSAGE REJECT and enters the Status phase reporting CHECK CONDITION with the sense key set to COMMAND ABORTED.

## BUS DEVICE RESET (0Ch)

---

**Caution** Due to the catastrophic nature of this message, it should be used cautiously in a multi-initiator system.

---

An initiator sends BUS DEVICE RESET to direct the drive to clear all current commands. When it recognizes this message, the drive aborts the command currently being executed and proceeds to the Bus Free state, writing any buffered data to tape.

A BUS DEVICE RESET message only resets the drive rather than all the devices on the bus, and so should be used to reset the drive instead of a hard reset.

## IDENTIFY (80h–FFh)

Either initiator or target sends these messages to establish the physical connection path between an initiator and target for a particular logical unit. The message byte can have the following bits set:

**Bit 7:** This bit is always set to one to distinguish these messages from other messages.

**Bit 6:** This bit can only be set to one by the initiator, and indicates that the initiator has the ability to disconnect and reconnect.

**Bits 5:** LUNTAB—this bit is always zero to indicate that the LUNTRN field (bits 2–0) is valid.

**Bits 4–3:** Reserved.

**Bits 2–0:** LUNTRN—these bits specify a Logical Unit Number in a target, and must always be set to zero because the HP C1533A is a single target, single logical unit device on the SCSI bus. The Logical Unit Number for the drive is fixed and is set to zero. Any other value causes the drive to enter the Status phase reporting CHECK CONDITION with the sense key set to COMMAND ABORTED.

It follows that the HP C1533A will only support the values 80 hex and C0 hex for the IDENTIFY message.

When sent from a target to an initiator during reconnection, an implied RESTORE POINTERS message is performed by the initiator before completing this message, since a RESTORE POINTERS is not usually sent.

---

## Status

A Status byte is sent from the drive to the host during the Status phase at the end of each command as specified in the SCSI specification, unless the command has been cleared by an ABORT message, by a BUS DEVICE RESET message, or by a hard reset.

The Status bytes that the HP C1533A returns are as follows:

**00h GOOD:** This status indicates that the drive has successfully completed the command.

**02h CHECK CONDITION:** Any error, exception, or abnormal condition that causes sense data to be set returns CHECK CONDITION. The REQUEST SENSE

command should be sent following this status to determine the nature of the error.

**08h BUSY:** The drive is busy. This status is returned whenever the drive cannot accept a command from an otherwise acceptable initiator. For example, if the HP C1533A is executing a command from an initiator which will not permit disconnect, and receives a command from a different initiator, The HP C1533A will return BUSY status. See "Command Queuing" in the Interface Implementation section.

Note that if an initiator which permits disconnect sends overlapped commands to the drive (in other words if it sends a new command while the previous one is still executing and the drive is disconnected), the first command will be aborted, and a CHECK CONDITION will be returned to the initiator to indicate that the command was aborted. The second command will *not* be executed.

---

**Note** The BUSY status returned by the drive must not be confused with DRIVE NOT READY. DRIVE NOT READY is returned as part of the Sense data following a REQUEST SENSE command and indicates that the drive is offline. The drive is regarded as offline if the tape is currently unloaded. Normally the tape will be ejected from the drive, but the drive is still offline if the tape is retained within the drive following a PREVENT MEDIA REMOVAL command issued by the host before the unload.

In the offline state, the host cannot perform any operation which would cause tape motion, in other words, write, read, verify and space commands. These commands will return a CHECK CONDITION status with a DRIVE NOT READY sense key set. The host may execute any diagnostic commands which do not access the tape, and usually a GOOD status will be returned. These commands are as follows:

- INQUIRY
- LOG SELECT
- LOG SENSE
- MODE SELECT
- MODE SENSE
- READ BLOCK LIMITS
- PREVENT/ALLOW MEDIUM REMOVAL
- READ BUFFER
- RECEIVE DIAGNOSTIC RESULTS
- REQUEST SENSE
- RESERVE/RELEASE UNIT
- SEND DIAGNOSTIC
- WRITE BUFFER

The host may load the tape when the unit is offline so long as the tape has been prevented from being ejected by using the PREVENT MEDIA REMOVAL command.

The TEST UNIT READY command is used to determine the online/offline status of the drive and so will return a CHECK CONDITION status with a DRIVE NOT READY sense key set if no tape is currently loaded or if a tape has just been loaded and the drive is now online.

- 
- 10h INTERMEDIATE GOOD:** This status is returned for every command in a series of linked commands except the last command, unless an error, exception or abnormal condition causes a CHECK CONDITION status or RESERVATION CONFLICT status to be set. If INTERMEDIATE GOOD is not returned, the chain of linked commands is broken; no further commands in the series are executed.
- 18h RESERVATION CONFLICT:** This status is returned by the drive if the host or another SCSI device attempts to access the drive when it has been reserved with the RESERVE UNIT command.

This chapter describes all HP C1533A SCSI commands. Parts of the chapter are based on sections 7 and 9 of the SCSI X3T9/89/042 draft revision specification.

## Summary

There are different groups of SCSI commands, each with a required length. Most of the commands are Group 0 commands which are 6 bytes long. Group 1 and Group 2 commands are 10 bytes long.

The following table is a summary of the SCSI commands for sequential access devices, showing the operation code and whether the command is supported:

Opcode			Opcode		
	(hex)	Type Command Name	(hex)	Type Command Name	
Group 0	00	S TEST UNIT READY	Group 1	2B	S LOCATE
	01	S REWIND		34	S READ POSITION
	02	S REQUEST BLOCK ADDRESS		39	U COMPARE
	03	S REQUEST SENSE		3A	U COPY AND VERIFY
	05	S READ BLOCK LIMITS		3B	S WRITE BUFFER
	08	S READ		3C	S READ BUFFER
	0A	S WRITE			
	0C	S SEEK BLOCK			
	0F	U READ REVERSE			
	10	S WRITE FILEMARKS			
	11	S SPACE			
	12	S INQUIRY	Group 2	40	U CHANGE DEFINITION
	13	S VERIFY		4C	S LOG SELECT
	14	U RECOVER BUFFERED DATA		4D	S LOG SENSE
	15	S MODE SELECT			
	16	S RESERVE UNIT			
	17	S RELEASE UNIT			
	18	U COPY			
19	S ERASE				
1A	S MODE SENSE				
1B	S LOAD/UNLOAD				
1C	S RECEIVE DIAGNOSTIC RESULTS				
1D	S SEND DIAGNOSTIC				
1E	S PREVENT/ALLOW MEDIUM REMOVAL				

**Key** S = Command is supported. See the command description for details.  
 U = Command is not supported.

## Command Details

The command descriptions in this section are listed in alphabetical order of command name. Each command is described briefly. This is followed by a list of pre-execution checks which are described below. The Command Descriptor Block (CDB) is then given, with details of the various parameter bits and fields.

In all Command Descriptor Blocks, bits 7-5 of byte 1 are the Logical Unit Number (LUN) field, and should be cleared to zero as recommended in the SCSI-2 standard.

In order to be SCSI-2 compatible, the following scheme is used to check the LUN:

SCSI-2 states that the LUN addressed by the host should appear in the IDENTIFY message sent to the target on selection. The HP C1533A checks the LUN field in the Command Descriptor Block for SCSI-1 and SCSI-2. However, SCSI-1 and SCSI-2 devices are allowed to coexist on the same bus, and SCSI-1 hosts will not send an IDENTIFY message. If the HP C1533A does not receive an IDENTIFY message on selection, it will check the LUN field in the Command Descriptor Block and expect it to be zero. If it is not zero, the drive will return a CHECK CONDITION status with an ILLEGAL REQUEST sense key set.

---

## Pre-Execution Checks

---

**Note** In compliance with the SCSI specification, the HP C1533A terminates a command with a CHECK CONDITION status and sets the sense key to ILLEGAL REQUEST when a reserved bit, byte, field or code is received which is not zero.

---

Before executing a command, the HP C1533A makes a number of checks. They fall into three categories:

- Checks on the command sent by the host. These ensure that no reserved or fixed fields have been set to illegal values. They check the syntax of commands, in other words the cross dependency of fields. For example, the Flag bit must not be set if the Link bit is not set.
- Checks to ensure that there is no asynchronous status pending for the host that sent the command. This status would be either a UNIT ATTENTION condition or a DEFERRED ERROR.
- Checks on media access abilities. These are performed for commands requiring access to the tape. A command is rejected if it attempts to access the tape when no tape is present or the tape is unloaded.

The checks are described below in alphabetical order. The usual order of execution is Illegal Field, Fixed Bit, Flag Link, Bad LUN, Reservation, Deferred Error, Unit Attention, Media Access, Media Write, Diagnostic Status, Parameter List.

### Bad LUN Check

This checks that the logical unit number specified in the lower three bits of the IDENTIFY message sent by the host was zero. This check is made after the logical unit number specified in the command descriptor has been seen to be zero.

- If no IDENTIFY message is supplied, the LUN is taken from the host's Command Descriptor Block.
- If the check fails, the host has sent conflicting logical unit identification information, and has requested use of a non-zero logical unit.
- If the LUN is not supported, and the host command is not REQUEST SENSE, CHECK CONDITION is reported to the host with a sense key of ILLEGAL REQUEST and additional sense of 2500h (logical unit not supported).
- If the LUN is not supported, and the host command is REQUEST SENSE then the original sense data is replaced with a sense key of ILLEGAL REQUEST and additional sense of 2500h (logical unit not supported), and this sense data is returned to the host immediately. After the command has completed successfully, the sense data is cleared.

### Deferred Error Check

A deferred error is generated when a command with immediate report fails after the report has been returned. The check looks to see if a deferred error exists for the host which sent the command, in other words, a deferred error for which CHECK CONDITION status has not yet been reported. If such an error exists, then the HP C1533A reports CHECK CONDITION. The sense data for the command is set to DEFERRED ERROR (which was generated when some previous command failed).

Note that if a UNIT ATTENTION condition and a DEFERRED ERROR condition both exist for an initiator, the DEFERRED ERROR condition will be reported first. This is because the operation leading to the deferred error must have been older than that leading to the unit attention. The HP C1533A reports the conditions in the order in which they arose.

### Diagnostic Status Check

This ensures that the drive is in a fit state to access the media. It does this by checking that there is no DIAGNOSTIC FAIL status within the drive. If the drive has failed diagnostics, CHECK CONDITION is reported with a sense key of HARDWARE ERROR and additional sense of 0401h (diagnostic failure on the controller board), or 0402h (diagnostic failure of the mechanism).

### Fixed Bit Check

For the READ, VERIFY and WRITE commands, a Fixed bit set to 1 indicates that the length parameter of the command is for fixed block mode. If fixed block mode is selected then the block size in the Mode Select block descriptor must not be zero. Otherwise CHECK CONDITION is reported with a sense key of ILLEGAL REQUEST and additional sense of 2400h (invalid field in CDB).

### Flag Link Check

This check ensures that the host has not set the Flag bit in the control byte of the command without setting the Link bit as well. If the test fails then CHECK CONDITION is reported with a sense key of ILLEGAL REQUEST and additional sense of 2400h (invalid field in CDB). The Flag field is identified as the bad field.

### Illegal Field Check

This checks that the host has set legal values in each of the fields, as described in the descriptions of each command that follow. If a field has been set to an illegal value, CHECK CONDITION is reported to the host with a sense key of ILLEGAL REQUEST and additional sense of 2400h (invalid field in CDB).

The drive scans the data in the Command Description Block from "left" (bit 7) to "right", and "down" (from byte 0 to byte n). It sets the field pointers to the first bit of the first bad field encountered. If the bad field is contained in a contiguous group of fixed fields, the pointers indicate the first bit of the first field in the group, even though the error may be in a later field in the group.

### Media Access Check

This checks if the drive is able to perform media access commands. If the media is inaccessible then CHECK CONDITION status is reported with a sense key of NOT READY. The additional sense will be set to one of the following:

- 3A00h (media not present). There is no cartridge in the drive.
- 0402h (initializing command required). A cartridge is present but is unloaded. The host must send a LOAD command before the tape can be accessed.
- 0400h (logical unit not ready, cause not reportable). A cartridge is being unloaded

or ejected. This sense is transient; when the unload or eject fails, it changes to 3A00h or 0402h.

- 0401h (logical unit in process of becoming ready). The HP C1533A is in the process of loading media and should be prepared to accept media access commands in future.

---

**Note** This means the drive is executing a LOAD command with immediate report on. The load could have been initiated by a host or it could be an autoloader. The host must effectively poll the HP C1533A by repeating the command until the media is loaded, at which point the command can be executed. If polling is not acceptable, it is advisable to disable autoloader. The host must then initiate all load operations, which it should do with IMMEDIATE REPORT disabled so that the HP C1533A reports status only when the load is complete.

---

- 3E00h (logical unit has not self-configured yet). The HP C1533A is executing its power-on self-test. When this is complete, the additional sense will change to 3A00h, 0401h, or 0402h depending on whether a cartridge is present at power-on and whether the drive is configured to autoloader the tape.

### Media Information Check

During power-on, knowledge of the whereabouts of the tape is unavailable. It is not possible to execute commands which assume that this knowledge is available until the drive has recovered from the power-on or reset. The test checks whether the drive knows if a cartridge is physically present in the drive.

If the information is not available, the test fails with CHECK CONDITION, a sense key of NOT READY, and additional sense of 3E00 (logical unit has not self-configured yet).

### Media Write Check

This checks whether the media is write-protected. If it is, then CHECK CONDITION is reported with a sense key of DATA PROTECT and additional sense of 2700h (operator selected write-protect).

### Parameter List Check

For LOG SELECT, MODE SELECT and COPY, the HP C1533A reads parameter lists. These are described under the command names later in this section. Checks are performed to test the following:

- Fixed fields have not been modified. Fixed fields are indicated by a number in round brackets following the field name.
- A field has been set to an invalid value.
- The syntax of the page of parameters has been violated—for example, where a particular value in one field imposes limitations on the valid range for another field.

If a field has been set to an illegal value, CHECK CONDITION is reported to the host with a sense key of ILLEGAL REQUEST and additional sense of 2600h (invalid field in parameter list).

The drive scans the data in the Command Description Block from “left” (bit 7) to “right”, and “down” (from byte 0 to byte n). It sets the field pointers to the first bit of the first bad field encountered. If the bad field is contained in a contiguous group of fixed fields, the pointers indicate the first bit of the first field in the group, even though the error may be in a later field in the group.



### Reservation Check

This checks to see if the HP C1533A has been reserved for use by a host, and if it has, whether the host is the same host that sent the command being executed.

If the drive has been reserved for some other host then RESERVATION CONFLICT status is reported. See the RESERVE UNIT and RELEASE UNIT commands.

### Unit Attention Check

This checks if a UNIT ATTENTION condition exists for the host which sent the command. If it does, the HP C1533A reports CHECK CONDITION status with a sense key of UNIT ATTENTION. The remaining sense data will be set according to the unit attention condition which exists. See Unit Attention Sense in the description of the REQUEST SENSE command which follows.

# ERASE

19h

## Description:

ERASE causes part or all of the remaining data within a partition to be erased, beginning at the current logical position. Write-behind data is written to tape before the ERASE is executed. Once an erase has started, it will be completed even if the SCSI operation is aborted.

## Pre-Execution Checks:

Illegal Field	Flag Link	Bad LUN	Reservation	Deferred Error
Unit Attention	Media Access	Media Write	Diagnostic Status	

## Command Descriptor Block:

		Bit								
		7	6	5	4	3	2	1	0	
Byte	0	Operation Code (19h)								
	1	Logical Unit Number (0)				Reserved (0)			Immed	Long
	2	Reserved (0)								
	3	Reserved (0)								
	4	Reserved (0)								
	5	Reserved (0)							Flag	Link

## CDB Fields:

Immed	0	The target will not return status until the ERASE has completed.
	1	Status will be returned immediately after the pre-execution checks have been completed.
Long	The Long bit controls the distance to be erased.	
	0	25 frames of amble are written. This is supported for compatibility reasons only, and serves no useful function; no useful space is created on the tape. It only causes the drive to streamfail. The logical tape position is not affected. A CHECK CONDITION for Early Warning EOM will only be given if the tape was logically positioned past EOT at the moment the command was issued.
	1	The drive writes EOD followed by ambles to the end of the partition (EOP).

## Erase Specific Status:

Event	Status	Key	Additional Sense
The erase fails	CHECK CONDITION	HARDWARE ERROR	5100h (erase failure)

# INQUIRY

12h

## Description:

INQUIRY tells the drive to return information about the drive parameters to the host.

## Pre-Execution Checks:

Illegal Field      Flag Link

**Note** The command is immune from most of the pre-execution checks other commands must pass, so that, for example, it can be executed while the drive is reserved for another host. Unit attention and deferred error conditions will be preserved and reported on subsequent commands.

## Command Descriptor Block:

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Operation Code (12h)							
	1	Logical Unit Number (0)				Reserved (0)			EVPD
	2	Page Code							
	3	Reserved (0)							
	4	Allocation Length							
	5	Reserved (0)						Flag	Link

## CDB Fields:

<b>Logical Unit Number</b>	Ignored. However, the LUN in a host's message-out IDENTIFY can affect the Inquiry data returned.
<b>EVPD</b>	<i>Enable Vital Product Data</i> 0 Normal inquiry data is returned. 1 A page of vital product data is returned.
<b>Page Code</b>	If the EVPD bit is zero the Page Code field must be zero. If the EVPD bit is set, The HP C1533A returns the page of inquiry data specified by the Page Code field: 00h Supported Vital Product Pages Page 01h ASCII Information Page 1 02h ASCII Information Page 2 03h ASCII Information Page 3 C0h Vendor-Unique Firmware Revision Page C1h Vendor-Unique Servo Revision Page
<b>Allocation Length</b>	0 An Allocation Length of zero means that no Inquiry data is to be transferred. >0 This specifies the maximum number of bytes that the HP C1533A will return. The drive will terminate the Data In phase when Allocation Length bytes have been transferred, or when all available Inquiry data has been transferred to the host, whichever is less.

## Returned Data:

INQUIRY returns its standard data if the EVPD bit is zero, or returns a page of data as specified by the Page Code field.

The following table shows the standard inquiry data format:

Commands  
INQUIRY

Standard Inquiry Data Format

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Peripheral Qualifier (0 or 3)				Peripheral Device Type (1)			
	1	RMB (1)	MChgr (0)	Device-Type Qualifier (0)					
	2	ISO Version (0)			ECMA Version (0)		ANSI-Approved Version (2)		
	3	AENC (0)	Reserved (0)			Response Data Format (2)			
	4	Additional Length (26h)							
	5	Reserved (0)							
	6	Reserved (0)							
	7	RelAdr (0)	WB32 (0)	WB16 (0)	Sync (1)	Linked (1)	Rsvd (0)	CmdQ (0)	SttRe (0)
	8	(MSB) Vendor Identification							
	15	("HP ") (LSB)							
	16	(MSB) Product Identification							
	31	("C1533A ") (LSB)							
	32	(MSB) Product Revision Level							
	35	(LSB)							
	36	(MSB) Manufacturing Date Code (HP Vendor-Unique)							
	39	("YY WW") 1960 + YY = Absolute Year, WW = week 1-52 (LSB)							
	40	Reserved (0)							
	41	Reserved (0)							
	42	HPCS Version (2)							

Inquiry Data Fields:

Peripheral Qualifier	0	A device exists at the LUN
	3	There is no device at the LUN.
Peripheral Device Type		If the LUN is zero, the drive will return a Peripheral Device Type of 01h (sequential-access device). If the LUN is not zero, the drive returns 1Fh.
RMB	1	The Removable Medium bit is one, indicating that the tape can be removed.
MChgr	0	The drive is not embedded in an autoloader.
Device-Type Modifier	0	This is a six-bit user defined code, set to zero.
ISO Version	0	This field is zero, indicating that the drive does not necessarily comply with the ISO version of SCSI.
ECMA Version	0	This field is zero, indicating that the drive does not necessarily comply with the ECMA version of SCSI.
ANSI Approved Version	2	This field is set to 2, indicating that the drive complies with the ANSI version of SCSI-2.
Response Data Format	2	This field has the value 2, indicating that the Inquiry Data format complies with the ANSI version of SCSI-2.
Additional Length	26h	The additional length field specifies the length in bytes of the rest of the Inquiry data, in this case, 38 (26h) bytes.
RelAdr	0	Relative Addressing is not supported.
WB32	0	32-bit wide data transfers are not supported. Only 8-bit transfers are supported.
WB16	0	16-bit wide data transfers are not supported.
Sync	1	Synchronous data transfer is supported
Linked	1	The drive supports linked commands, so this bit is set to one.
CmdQ	0	Command Queuing is not supported. This means that the drive will not queue multiple commands from a single host. It will, however, queue single commands from several hosts.

<b>StRe</b>	0	The drive responds to the reset condition with a hard reset, so the Soft Reset bit is never set.
<b>Vendor Identification</b>		This field contains eight bytes of ASCII data identifying the vendor of the product as 'HP'. This and the next two fields are left-aligned with the unused bytes at the end of the fields filled with space characters (20h).
<b>Product Identification</b>		This field contains sixteen bytes of ASCII data identifying the product as 'C1533A'.
<b>Product Revision Level</b>		This field contains four bytes of ASCII data, which define the revision level of the Controller firmware. Note that in order to read a precise version of the firmware, a host should read the vendor-unique Firmware Revisions page (C0h) described on the next page instead of this field.
<b>Manufacturing Date Code</b>		This field contains four bytes of Vendor-Unique data which identify variants of the firmware for particular customers.
<b>HPCS version</b>		This is a single byte Vendor-Unique field identifying which version of HPCS the drive supports.

## Inquiry Data Pages

The following diagram shows the Supported Pages Page:

Supported Vital Product Data Pages

		<i>Bit</i>							
		7	6	5	4	3	2	1	0
<i>Byte</i>	0	Peripheral Qualifier (0)				Peripheral Device Type (1)			
	1	Page Code (00h)							
	2	Reserved (0)							
	3	Page Length (06h)							
	4	Supported Pages Code (00h)							
	5	ASCII Information Page 1 (01h)							
	6	ASCII Information Page 2 (02h)							
	7	ASCII Information Page 3 (03h)							
	8	Vendor-Unique Firmware Revision Page (C0h)							
	9	Vendor-Unique Servo Revision Page (C1h)							

The ASCII Information Pages:

The three ASCII Information Pages are as follows:

ASCII Information Page 1

		<i>Bit</i>							
		7	6	5	4	3	2	1	0
<i>Byte</i>	0	Peripheral Qualifier (0)				Peripheral Device Type (1)			
	1	Page Code (01h)							
	2	Reserved (0)							
	3	Page Length (19h)							
	4	ASCII Length (18h)							
	5	MSB	"FRU 1 fail. Brain Dead."						LSB
	28								

Commands  
INQUIRY

ASCII Information Page 2

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Peripheral Qualifier (0)				Peripheral Device Type (1)			
	1	Page Code (02h)							
	2	Reserved (0)							
	3	Page Length (19h)							
	4	ASCII Length (18h)							
	5	MSB	"FRU 2 fail. Brain Dead."						LSB
28									

ASCII Information Page 3

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Peripheral Qualifier (0)				Peripheral Device Type (1)			
	1	Page Code (03h)							
	2	Reserved (0)							
	3	Page Length (19h)							
	4	ASCII Length (18h)							
	5	MSB	"FRU 3 fail. Brain Dead."						LSB
28									

Vendor-Unique Information Pages:

The Vendor-Unique Information Pages are as follows:

Vendor-Unique Firmware Revisions

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Peripheral Qualifier (0)				Peripheral Device Type (1)			
	1	Page Code (C0h)							
	2	Reserved (0)							
	3	Page Length (89h)							
	4	ASCII Length (88h)							
	5	MSB	5-29 "Firmware Rev = xxxxxxxxxxxxxx"						LSB
:		30-72 "Build Date = ddd mmm dd yy. hh:mm GMT" (e.g. "Fri Nov 23 90. 20:09 GMT")							
140		72-140 "Copyright 1990 by Hewlett-Packard Company. All rights are reserved."							

Vendor-Unique Servo Revisions

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Peripheral Qualifier (0)				Peripheral Device Type (1)			
	1	Page Code (C1h)							
	2	Reserved (0)							
	3	Page Length (15h)							
	4	ASCII Length (14h)							
	5	MSB	"Servo Rev = RRR.VVV"						LSB
:		(RRR = Revision, VVV = Version)							
24									

# LOAD/UNLOAD

1Bh

## Description:

LOAD/UNLOAD tells the target to load or unload a tape.

## Pre-Execution Checks:

Illegal Field	Flag Link	Bad LUN	Reservation
Deferred Error	Unit Attention	Diagnostic Status ( <i>load only</i> )	

## Command Descriptor Block:

		Bit								
		7	6	5	4	3	2	1	0	
Byte	0	Operation Code (1Bh)								
	1	Logical Unit Number (0)				Reserved (0)				Immed
	2	Reserved (0)								
	3	Reserved (0)								
	4	Reserved (0)					EOT (0)	ReTen (0)	Load	
	5	Reserved (0)						Flag	Link	

## CDB Fields:

Immed	0	If this bit is zero, status is returned after the operation has completed.
	1	When this bit is set, status is returned as soon as the load or unload operation is started.
EOT	This field is ignored.	
ReTen	This field is ignored.	
Load	0	If the bit is zero, the drive performs an <i>unload</i> operation; it writes any buffered data to the tape, rewinds the tape to BOM and ejects it.
	1	When this bit is set, the drive performs a <i>load</i> operation; it loads the tape and positions it at the beginning of Partition 0.

## Load/Unload Specific Status:

Event	Status	Key	Additional Sense
Cannot read media	CHECK CONDITION	MEDIA ERROR	3002h (Cannot read media, incompatible format)
Format corrupted	CHECK CONDITION	MEDIA ERROR	3100h (Media format corrupted)
No tape present	CHECK CONDITION	NOT READY	3A00h (medium not present)
Tape being loaded	CHECK CONDITION	NOT READY	0401h (LUN in process of becoming ready)
Tape being unloaded or ejected	CHECK CONDITION	NOT READY	0400h (LUN not ready, cause unreportable)
After successful load	GOOD	UNIT ATTENTION posted at all other initiators*	2800h (not ready to transition)
When UNIT ATTENTION cleared	UNSOLICITED POSITIONAL SENSE	NO SENSE	0004h (BOP detected), EOM bit set

\* An exception to this is when a load occurs when the tape is already loaded and positioned within partition 0. In this case, UNIT ATTENTION sense is not generated for other initiators.

### Loading a Tape:

- If the tape is already loaded, it is positioned at BOP for partition 0. (For a tape already positioned in partition 0, this is equivalent to responding to a REWIND command.) The system area is not updated, so the "current" fields of the tape log are not reset.
- After a successful load, the tape is logically positioned at the beginning of partition 0. Media access commands are then permissible.
- If a tape is in the process of being loaded, the drive waits for the load operation to complete.
- Status is generated as follows:
  - If the Load command was issued by the current host, and completed successfully, GOOD status is returned to the host. No further action is taken.
  - If the Load command was issued by a different host and completed successfully, CHECK CONDITION status is returned to the current host, with a sense key of UNIT ATTENTION and additional sense of 2800h (not ready to transition).
  - If the Load command was issued by the current host and failed, CHECK CONDITION is reported, since a deferred error condition will exist for the current host.
  - If the load was issued by a different host and failed, the load will be re-attempted for the current host.

---

**Note** The drive cannot detect the presence of media during power-on self-tests. If a Load command is received during this period, CHECK CONDITION is returned to the host. The sense key is set to NOT READY and additional sense is set to 3E00h (LUN has not self-configured yet).

---

### Unformatted Tape

If an unformatted tape is inserted, the drive loads it and behaves as though the tape has a single partition.

### Unloading a Tape

- Unload can also be initiated by pressing the Unload button, unless PREVENT MEDIA REMOVAL is in effect.
- All media access commands are rejected after an Unload command.
- If an Unload command is received while unloading is already in progress, the command will be queued until the unload completes. If the unloading operation is successful, GOOD status is reported. If not, another unload operation is initiated for the host that issued the command.
- If an Unload command is received while there is no cartridge present, or a cartridge is present-but-not-loaded and medium removal is prevented, GOOD status is reported.

The unload procedure is as follows:

- 1 Any buffered data is written to tape.
- 2 The system area of the current partition is updated.
- 3 The tape is rewound to BOM.
- 4 If media removal is not prevented, the cartridge is ejected. If media removal is prevented, the tape is retained in the drive and a Load command can load the tape again.



# LOCATE

2Bh

## Description:

The LOCATE command moves the current logical position to a new position specified by the following:

- 1 A new partition
- 2 The offset from the start of the new partition in terms either of blocks (ignoring marks), or the count of all blocks and marks. The BT (Block address Type) bit selects which. Block 0 is located at the start of the partition.

Any unwritten data is written to tape before the command is executed.

## Pre-Execution Checks:

Illegal Field Deferred Error	Flag Link Unit Attention	Bad LUN Media Access	Reservation Diagnostic Status
---------------------------------	-----------------------------	-------------------------	----------------------------------

## Command Descriptor Block:

		Bit							
		7	6	5	4	3	2	1	0
	0	Operation Code (2Bh)							
	1	Logical Unit Number (0)			Reserved (0)		BT	CP	Immed
	2	Reserved (0)							
Byte	3	(MSB)		Block Address					
	6								
	7	Reserved (0)							
	8	Partition							
	9	Reserved (0)						Flag	Link

## CDB Fields:

<b>BT</b>	<i>Block address Type</i>	0	The Block Address field is the total number of blocks, filemarks and setmarks from the beginning of the partition to the current logical position.
		1	The value in the Block Address field is the number of blocks from the beginning of the partition to the new position.
<b>CP</b>	<i>Change Partition</i>	0	No change of partition is made.
		1	The drive will change partitions to the partition specified in the Partition field.
<b>Immed</b>		0	The HP C1533A will not report status until the locate operation has completed.
		1	The drive flushes any unwritten data and, if necessary, switches partitions before reporting to the host. The drive then moves to the required position within the partition after reporting status.
<b>Block Address</b>	The Block Address specifies either the number of blocks or the total number of blocks plus marks from the beginning of the partition to the new position. The BT bit determines which measure is used. The Block Address must be greater than zero, otherwise CHECK CONDITION is reported.		
<b>Partition</b>	If the CP bit is set, the Partition field specifies which partition to change to. 0 and 1 are the only valid values for a 2-Partition tape. The field must be 0 for 1-Partition tape. If the CP bit is not set, the value of the field is ignored.		

Commands  
LOCATE

Locate Specific Status:

Event	Status	Key	Additional Sense
EOD encountered	CHECK CNDTN	BLANK CHECK	0005h (EOD encountered). EOM bit set.
EOP encountered	CHECK CNDTN	MEDIUM ERROR	0002h (EOP encountered). EOM bit set.
Failed to read data—media error or non-fatal drive error	CHECK CNDTN	MEDIUM ERROR	3B00h (sequential positioning error)
EWEOM encountered*	CHECK CNDTN	NO SENSE	0002h (EOP/M detected)
Data format corrupt	CHECK CNDTN	MEDIUM ERROR	3001h (can't read media, unknown format)
Cannot decompress the entity in which the record corresponding to the target block address lies	CHECK CNDTN	MEDIUM ERROR	3B00h (sequential positioning error)

\* Early Warning End of Medium (EWEOM) is only reported if the LOCATE operation is successful and the REW bit of the DEVICE CONFIGURATION mode page is set.

# LOG SELECT

4Ch

## Description:

LOG SELECT allows the host to clear logs in the drive. The logs are maintained by the drive, and contain statistical information about the drive and the media. Since access to individual parameters within log pages is not supported, the host is restricted to resetting complete log pages only. For information about log page format, parameters and supported pages, see the LOG SENSE command description.

**Note** Logs cannot be saved since the drive does not possess non-volatile RAM.

## Pre-Execution Checks:

Illegal Field    Flag Link    Bad LUN    Reservation    Deferred Error    Unit Attention    Diagnostic Status

## Command Descriptor Block:

		Bit								
		7	6	5	4	3	2	1	0	
Byte	0	Operation Code (4Ch)								
	1	Logical Unit Number (0)			Reserved (0)			PCR	Rsvd (0)	
	2	PC		Reserved (0)						
	3	Reserved (0)								
	6	Reserved (0)								
	7	(MSB)	Parameter List Length						(LSB)	
	8	Reserved (0)								
	9	Reserved (0)						Flag	Link	

## CDB Fields:

<b>PCR</b>	<i>Parameter Code Reset</i>
	0 The function performed is defined by the PC field. The Parameter List Length must be greater than zero.
	1 All log pages will be cleared which can be cleared. The Parameter List Length must be zero.
<b>PC</b>	The <i>Page Control</i> field defines the type of parameter values to be selected:
	00b Current Threshold Values—NO-OP is performed
	01b Current Cumulative Values—logs are cleared as specified by the parameter data
	10b Default Threshold Values—NO-OP is performed
	11b Default Cumulative Values—all internal logs are cleared
<b>Parameter List Length</b>	0 No data is to be transferred. If the PCR bit is 1, all log pages which can be cleared are cleared.
	>0 Specifies the length in bytes of the LOG SELECT parameter list to be transferred from the initiator to the target during the Data-Out phase. The value of this field must be a multiple of 4, and less than 4 × number of pages. The PC field must be 01b.

Since the host can only reset complete log pages, it must send the log page header of the page to be cleared with the Page Length field set to 0 during the DATA OUT phase. The following pages can be cleared in this manner:

<u>Page Code</u>	<u>Page Description</u>
02h	Write Error Counter Page
03h	Read Error Counter Page
31h	Tape Capacity Log Page
32h	Buffer Trace Log Page
33h	Device Trace Log Page
34h	Write Frames Error Counters Page
35h	Read Frames Error Counters Page
36h	Bad Group Log Page
37h	Drive Counters Page
38h	Mechanism Counters Page
39h	Data Compression Transfer Log Page
3Ah	Data Compression Trace Log Page
3Bh	Channel Trace Log Page

Because of this method of clearing the log pages, the Parameter List Length field must be a multiple of 4, the length of the Log Page header. Otherwise the command will be terminated with a CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and an additional sense code of INVALID FIELD IN CDB.

---

**Note** The Tape log and Tape Capacity page cannot be cleared.

---

# LOG SENSE

4Dh

## Description:

LOG SENSE allows the host to read logs from the HP C1533A. A single log is returned with each invocation of LOG SENSE.

## Pre-Execution Checks:

Illegal Field    Flag Link    Bad LUN    Reservation    Deferred Error    Unit Attention

## Command Descriptor Block:

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Operation Code (4Dh)							
	1	Logical Unit Number (0)				Reserved (0)			
	2	PC		Page Code					
	3	Reserved (0)							
	6	Reserved (0)							
	7	(MSB)		Allocation Length				(LSB)	
	8	Reserved (0)							
	9	Reserved (0)						Flag	Link

## CDB Fields:

<b>PC</b>	<p>The <i>Page Control</i> field defines the type of parameter values to be selected:</p> <p>00b <b>Current Threshold Values</b>—cannot be changed. These are the maximum values that each parameter can attain. Note that for some parameters the term threshold value has no meaning—see the descriptions of individual pages and parameters which follow for more details.</p> <p>01b <b>Current Cumulative Values</b>—the values computed since the last reset of the drive (either by power-cycling, BUS DEVICE RESET or SCSI RESET). Note that for some parameters these values cannot be reset by any method—see the descriptions of individual pages and parameters which follow for more details.</p> <p>10b <b>Default Threshold Values</b>—same as the Current Threshold Values</p> <p>11b <b>Default Cumulative Values</b>—the values to which each parameter is initialized when reset, as described above. See the descriptions of individual pages and parameters which follow for more details.</p> <p>The PC field has no effect on the data returned when the selected log contains event or trace codes rather than counts.</p>
-----------	--

Commands  
LOG SENSE

---

<b>Page Code</b>	The <i>Page Code</i> field identifies which log page is being requested by the host, as follows:
	00h List of Supported Pages Page
	02h Write Error Counter Page
	03h Read Error Counter Page
	30h Tape Log Page
	31h Tape Capacity Log Page
	32h Buffer Trace Log Page
	33h Device Trace Log Page
	34h Write Frames Error Counters Page
	35h Read Frames Error Counters Page
	36h Bad Group Log Page
	37h Drive Counters Page
	38h Mechanism Counters Page
	39h Data Compression Log Page
	3Ah Data Compression Trace Log Page
	3Bh Channel Trace Log Page

---

<b>Allocation Length</b>	The Allocation Length field specifies the maximum number of bytes of data which should be returned to the host. The drive will return the entire log or Allocation Length bytes, whichever is the lesser.
--------------------------	---

## Log Parameter Pages

### Log Page Descriptor

Each page begins with a 4-byte page header followed by one or more variable-length parameter structures. LOG SENSE returns only the single log page specified in the Page Code field of the command descriptor block. With LOG SELECT however, several log page headers can be sent to clear several pages simultaneously.

### Supported Log Pages Page

The Supported Log Pages page has the following format:

		<i>Bit</i>							
		7	6	5	4	3	2	1	0
0		Page Code (00h)							
1		Reserved (0)							
2	(MSB)	Page Parameter Length (000Fh)							
3		(LSB)							
4		Supported Pages page code (00h)							
5		Write Error Counters page code (02h)							
6		Read Error Counters page code (03h)							
7		Tape Log page code (30h)							
8	<i>Byte</i>	Tape Capacity Log page code (31h)							
9		Buffer Trace Log page code (32h)							
10		Device Trace Log page code (33h)							
11		Write Frames Error Counters page code (34h)							
12		Read Frames Error Counters page code (35h)							
13		Bad Group Log page code (36h)							
14		Drive Counters page code (37h)							
15		Mechanism Counters page code (38h)							
16		Data Compression Log page code (39h)							
17		Data Compression Trace Log page code (3Ah)							
18		Channel Trace Log page code (3Bh)							

The PC field of the LOG SENSE command has no effect on the values returned in this page. The Supported Log Pages Page is not a valid page to send during LOG SELECT.

## Write and Read Error Counter Pages

The Write and Read Error Counters Pages have the following format. All fields in these logs are initialized to zero in the following circumstances:

- At power-on
- After a SCSI reset
- After a tape load
- Following configuration for Error-Rate Testing
- After a user sends an request for initialization through LOG SELECT

		Bit								
		7	6	5	4	3	2	1	0	
Byte	0	Page Code (02h Write Error Counters or 03h Read Error Counters)								
	1	Reserved (0)								
	2	(MSB)				Page Length (1Ch)				(LSB)
	3									
—Error Counter Descriptors—										
	0	(MSB)				Parameter Code				(LSB)
	1									
	2	DU (0)	DS (1)	TSD (0)	ET (0)	TMC (0)		Rsvd (0)	LP (0)	
	3	Parameter Length (m)								
	4	(MSB)				Parameter Value				(LSB)
	3+m									

### Write and Read Error Counters Page Fields:

**Note** The Current or Default Cumulative, or Current or Default Threshold values are returned according to the PC field of LOG SENSE.

The DU (Disable Update), DS (Disable Save), TSD (Target Save Disable), ET (Enable Threshold Comparison), TMC (Threshold Met Criteria) and LP (List Parameter) fields are collectively referred to as the *control byte*.

**Parameter Code** The Parameter Code field identifies which parameter of data is being transferred. The following parameter codes are supported for each Error Counters page. The parameter types for write and read operations are described below:

Param. Code	Description	Length (bytes)	Default Cumulative	Default Threshold
0003h	Total Errors Corrected (Soft Errors)	2	0	FF FFh
0004h	Total Times Correction Algorithm Processed	4	0	FF FF FF FFh
0005h	Total Groups Processed	4	0	FF FF FF FFh
0006h	Total Errors Uncorrected (Hard Errors)	2	0	FF FFh
<b>DS</b>	1	This tells the host that saveable parameters are disabled (i.e. not supported).		
<b>LP</b>	0	Must be zero; the parameter is a data counter rather than a list parameter.		
<b>DU, TSD, ET, TMC</b>	0	Must be zero. For a full description of these fields refer to the SCSI-2 standard.		

#### Write Parameters

**Total Errors Corrected** The number of groups written where write 'retries' were necessary; in other words, the whole group was rewritten. The count is only incremented once for a particular group, even if many retries occurred for that group.

**Total Times Correction Algorithm Processed** The number of frames which needed to be rewritten after the read-after-write check. It does not include the number of 'dummy' frames



rewritten between the original version of the frame and its read-after-write retry.

- Total Groups Processed** The total number of groups written.
- Total Errors Uncorrected** The number of groups which could not be written successfully after the maximum number of write retries had been attempted.

*Read Parameters*

- Total Errors Corrected** The number of groups which required a retry (involving a reposition of the tape) during a read. The count is only incremented once for a particular group, even if many retries occurred for that group. If the group read is part of an N-Group written sequence, the count is incremented by the number of groups within the sequence.
- Total Times Correction Algorithm Processed** The number of groups read where C1 and C2 correction was insufficient, and C3 error correction was needed. The count is incremented on each retry of a group.
- Total Groups Processed** The total number of groups read.
- Total Errors Uncorrected** The number of groups which could not be read after C3 correction and the maximum number of retries were performed.

## Tape Log Page

The Tape Log page is used to return information to the host in a format which conforms to the LOG SENSE command page format. The Tape Log cannot be cleared. If no tape is loaded, the information returned is undefined. If a blank tape is loaded, the information returned will indicate a load count of 1 and all other fields will be zero. The PC field of LOG SENSE is ignored for this page.

The page has the following format:

		Bit									
		7	6	5	4	3	2	1	0		
Byte	0	Page Code (30h)									
	1	Reserved (0)									
	2	(MSB)		Page Length (58h)							
	3									(LSB)	
—Tape Log Descriptors—											
Byte	0	(MSB)		Parameter Code							
	1									(LSB)	
	2	DU (0)	DS (1)	TSD (0)	ET (0)	TMC (0)		Rsvd (0)	LP (0)		
	3	Parameter Length (m)									
	4	(MSB)		Parameter Value							
	3+m									(LSB)	

### Tape Log Page Fields:

Supported Parameter Codes in this page and their meanings are as follows:			
Parameter Codes	Code	Description	Length (bytes)
	0001h	Current Number of Groups Written	3
	0002h	Current RAW Retries	2
	0003h	Current Number of Groups Read	3
	0004h	Current C3 ECC Retries	2
	0005h	Previous Number of Groups Written	3
	0006h	Previous RAW Retries	2
	0007h	Previous Number of Groups Read	3
	0008h	Previous C3 ECC Retries	2
	0009h	Total Number of Groups Written	4
	000Ah	Total RAW Retries	3
	000Bh	Total Number of Groups Read	4
	000Ch	Total C3 ECC Retries	3
	000Dh	Load Count	2

'Current' refers to the current load of the tape. 'Previous' refers to the last tape load. 'Total' refers to the life of the tape since it was first initialized, including the current load.

### Tape Partitions:

The information returned by the Tape Log is specific to the current partition.

<b>Partition 0</b>	The counts indicate read, write and ECC activity in this partition. For a single partition tape, this means the whole tape. The <i>load count</i> is a special case. For a single partition tape, this indicates the total number of loads (including the current load) since the tape was last formatted. For a 2-partition tape, the Partition 0 load count is set to zero.
<b>Partition 1</b>	The counts indicate the total read, write and ECC for both partitions. The load count indicates the total number of loads for both partitions (including the current load) since the tape was last formatted, with the count for Partition 0 set to zero.

### Use of the Tape Log Page

An estimate of reading and writing performance can be made by the ratio of 'RAW retries' to 'groups written', and 'C3 ECC retries' to 'groups read'. An indication of tape degradation can be gained by comparing these ratios for 'current', 'previous', and 'total' entries.

Any attempt to change this page using LOG SELECT will result in a CHECK CONDITION with the ILLEGAL REQUEST sense key value returned. The current Cumulative values will be returned regardless of the Page Control field in the LOG SENSE Command Descriptor Block.

### Tape Capacity Log Page

The Tape Capacity Log page is used to indicate the approximate capacity of the tape currently being used. The PC field of the LOG SENSE command is ignored for this field.

The 36 bytes of the log page have the following format:

		Bit								
		7	6	5	4	3	2	1	0	
Byte	0	Page Code (31h)								
	1	Reserved (0)								
	2	(MSB)		Page Length (20h)						(LSB)
	3									
—Capacities—										
Byte	0	(MSB)		Parameter Code						(LSB)
	1									
	2	DU (0)	DS (1)	TSD (0)	ET (0)	TMC (11)		BES (0)	LP (0)	
	3	Parameter Length (04h)								
	4	(MSB)		Parameter Value						(LSB)
	7									
	7									

### Tape Capacity Log Page Parameters:

Code	Value
0001	Remaining capacity, partition 0 (kilobytes).
0002	Remaining capacity, partition 1 (kilobytes). For a single partition tape, it will be zero.
0003	Maximum capacity, partition 0 (kilobytes)
0004	Maximum capacity, partition 1 (kilobytes). For a single partition tape, it will be zero.

The following points affect these values:

- The remaining capacity for the partition which is not current will be the same as the maximum capacity for that partition.
- Remaining capacity values are only valid after the successful completion of the following commands in non-immediate mode:

LOAD	LOCATE	MODE SELECT	READ	REWIND
SPACE	VERIFY	WRITE	WRITE FILEMARKS	

The values after any subsequent command cannot be relied on unless the command is a sense-type command which does not cause any tape motion.

- Maximum capacity values are only valid when the tape has completed a load sequence. If an immediate mode LOAD is made, LOG SENSE will not return valid values until the tape has been successfully loaded and tape movement has ceased.
- Capacities are given in kilobytes of user data. If data compression is used, the capacities are specified as though the drive is in "pass-through" mode. The data-compression factor is not considered.

## Accuracy of Capacity Estimates

Remaining capacities are the capacities from the current physical position on the tape, not the unused capacity of the tape. This is particularly important if the current physical position is within the used data area.

### *Allowances*

System areas and Read-After-Write retries can only be accommodated by estimates based on the history of patterns of tape usage from earlier drives. In the estimates used in the log, conservative figures are used to ensure that the capacity stated is actually available. It will usually be possible to write considerably more data than is stated.

- Read-After-Write Retries** An allowance is made of 5 frames per group (approximately 22%). The actual occurrence is likely to be far less than this.
- Early-Warning EOM Area** Approximately 10 megabytes are allowed for between the Early-Warning EOM point and the physical EOM. Approximately 1 megabyte is allowed for the Early-Warning EOM buffering factor. The assumption is that the buffer is full, and as a result reports Early-Warning up to 1 megabyte earlier than is necessary.
- EOD** Approximately 1.6 megabytes are allowed for writing EOD frames.
- BOP** If the tape is at BOM, approximately 1 megabyte is allowed for the System Area and Vendor Group.
- N-Group Writing** The number of frames available is divided by the N-Group factor to give the apparent capacity available to the user.

The remaining capacity is then calculated as follows:

$$\text{capacity} = \frac{\text{frames remaining in partition} - \text{allowances}}{N \text{ Group factor} \times \text{no. of frames per megabyte}} \text{ megabytes}$$

### *Physical Variations*

Variations in tape dimensions, capstan diameter, supply reel hub diameter, and supply reel velocity can all result in anomalies in the estimated capacities:

- Different loads of the same tape can return different maximum capacities of partition 0.
- At BOP, the remaining capacity can be less than the maximum capacity.

## Buffer Trace Log Page

The Buffer Trace log consists of  $n$  parameters, each representing an event which has occurred concerning the buffer code. The first entry is the oldest event, and the last the most recent.

$n$  parameters are always available even if the drive has just been powered up and the log is not filled. The oldest events are overwritten as new buffer events occur, so that logging never stops.

The PC field of the LOG SENSE command has no effect on the data returned in this log.

The log page has the following format:

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Page Code (32h)							
	1	Reserved (0)							
	2	(MSB)	Page Length						(LSB)
	3								
—Buffer Event Traces—									
	0	(MSB)	Parameter Code ( $n$ )						(LSB)
	1								
Byte	2	DU (0)	DS (1)	TSD (0)	ET (0)	TMC (00)	Rsvd (0)	LP (0)	
	3	Parameter Length (08h)							
	4	(MSB)	Trace Point ID						(LSB)
	5								
	6	(MSB)	Time Stamp						(LSB)
	7								
	8	(MSB)	Argument 1						(LSB)
	9								
	10	(MSB)	Argument 2						(LSB)
	11								

### Buffer Trace Log Page Fields:

<b>Trace Point ID</b>	This is a unique code identifying the event which occurred
<b>Time Stamp</b>	The value of the lower 16 bits of the system clock at the time the event occurred. The system clock increments once every 512 $\mu$ s.
<b>Argument 1</b>	A parameter associated with the event
<b>Argument 2</b>	A parameter associated with the event

## Device Trace Log Page

The Device Trace log consists of  $n$  parameters, each representing an event which has occurred concerning the device code. The first entry is the oldest event, and the last the most recent.

$n$  parameters are always available even if the drive has just been powered up and the log is not filled. The oldest events are overwritten as new device events occur, so that logging never stops.

The PC field of the LOG SENSE command has no effect on the data returned in this log.

The log page has the following format:

		Bit								
		7	6	5	4	3	2	1	0	
Byte	0	Page Code (33h)								
	1	Reserved (0)								
	2	(MSB)	Page Length						(LSB)	
	3									
—Device Event Traces—										
Byte	0	(MSB)	Parameter Code ( $n$ )						(LSB)	
	1									
	2	DU (0)	DS (1)	TSD (0)	ET (0)	TMC (00)		Rsvd (0)	LP (0)	
	3	Parameter Length (08h)								
	4	(MSB)	Trace Point ID						(LSB)	
	5									
	6	(MSB)	Time Stamp						(LSB)	
	7									
	8	(MSB)	Argument 1						(LSB)	
	9									
	10	(MSB)	Argument 2						(LSB)	
	11									

### Device Trace Log Page Fields:

<b>Trace Point ID</b>	This is a unique code identifying the event which occurred
<b>Time Stamp</b>	The value of the lower 16 bits of the system clock at the time the event occurred. The system clock increments once every 512 $\mu$ s.
<b>Argument 1</b>	A parameter associated with the event
<b>Argument 2</b>	A parameter associated with the event

## Write and Read Frame Error Counter Pages

The Write and Read Frame Error Counter Pages have the following format. All fields in these logs are initialized to zero in the following circumstances:

- At power-on
- After a SCSI reset
- After a tape load
- Following configuration for Error-Rate Testing
- After a user sends an request for initialization through LOG SELECT

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Page Code (34h Write Frame Error or 35h Read Frame Error Counters)							
	1	Reserved (0)							
	2	(MSB)	Page Length (38h/48h)						(LSB)
	3								
--Error Counter Descriptors--									
	0	(MSB)	Parameter Code						(LSB)
	1								
	2	DU (0)	DS (1)	TSD (0)	ET (0)	TMC (0)		Rsvd (0)	LP (0)
	3	Parameter Length (4h)							
	4	(MSB)	Parameter Value						(LSB)
	7								

## Write and Read Error Counters Page Fields:

**Note** The Current or Default Cumulative, or Current or Default Threshold values are returned according to the PC field of LOG SENSE.

The DU (Disable Update), DS (Disable Save), TSD (Target Save Disable), ET (Enable Threshold Comparison), TMC (Threshold Met Criteria) and LP (List Parameter) fields are collectively referred to as the *control byte*.

**Parameter Code** The Parameter Code field identifies which parameter of data is being transferred. The following parameter codes are supported for each Error Counters page. Positive and negative tracks mean tracks with +ve and -ve azimuth:

Para. Code	Description	Length (bytes)	Default Cumulative	Default Threshold
0001h	Frames Read or Written	4	0	FF FF FF FFh
0002h	Main Data C1 Block Errors—positive tracks	4	0	FF FF FF FFh
0003h	Main Data C1 Block Errors—negative tracks	4	0	FF FF FF FFh
0004h	Sub Area 0 C1 Block Errors—positive tracks	4	0	FF FF FF FFh
0005h	Sub Area 1 C1 Block Errors—positive tracks	4	0	FF FF FF FFh
0006h	Sub Area 0 C1 Block Errors—negative tracks	4	0	FF FF FF FFh
0007h	Sub Area 1 C1 Block Errors—negative tracks	4	0	FF FF FF FFh
0008h	Total Retry count (Read Frames log only)	4	0	FF FF FF FFh
0009h	Read C2 Uncorrectable block (Read Frames log only)	4	0	FF FF FF FFh
DS	1	This tells the host that saveable parameters are disabled (i.e. not supported)		
LP	0	Must be zero; the parameter is a data counter rather than a list parameter.		
DU, TSD, ET, TMC	0	Must be zero. For a full description of these fields refer to the SCSI-2 standard.		

## Bad Group Log Page

The Bad Group log consists of  $n$  parameters, each representing a group error event which has occurred concerning the device code. The first entry is the oldest event, and the last the most recent.

$n$  parameters are always available even if the drive has just been powered up and the log is not filled. The oldest events are overwritten as new channel events occur, so that logging never stops.

The PC field of the LOG SENSE command has no effect on the data returned in this log.

All fields in this log are initialized to zero in the following circumstances:

- At power-on
- After a SCSI reset
- After a tape load
- Following configuration for Error-Rate Testing
- After a user sends an request for initialization through LOG SELECT

The log page has the following format:

		Bit								
		7	6	5	4	3	2	1	0	
Byte	0	Page Code (36h)								
	1	Reserved (0)								
	2	(MSB) Page Length								
	3	(LSB)								
—Bad Group Event Traces—										
Byte	0	(MSB) Parameter Code ( $n$ )								
	1	(LSB)								
	2	DU (0)	DS (1)	TSD (0)	ET (0)	TMC (00)	Rsvd (0)	LP (0)		
	3	Parameter Length (03h)								
	4	Error Code								
	5	(MSB) Group Count								
	6	(LSB)								

### Error Codes:

The Error codes that can occur are as follows:

Code	Description
0x20	Group read after write retry limit exceeded
0x21	Write hard error
0x22	Retrying a group (writing)
0x23	Group read retry limit exceeded
0x24	Retrying a group (reading)
0x25	C2 uncorrectable tracks
0x26	Read drive error
0x27	Group needed C3 correction
0x28	Group needed Read-after-Write



## Drive Counters Log Page

The Drive Counters Log Page has the following format. All fields in this log are initialized to zero in the following circumstances:

- At power-on
- After a SCSI reset
- After a tape load
- Following configuration for Error-Rate Testing
- After a user sends an request for initialization through LOG SELECT

		Bit								
		7	6	5	4	3	2	1	0	
Byte	0	Page Code (37h)								
	1	Reserved (0)								
	2	(MSB)		Page Length (18h)						(LSB)
	3									
—Drive Counter Descriptors—										
	0	(MSB)		Parameter Code						(LSB)
	1									
	2	DU (0)	DS (1)	TSD (0)	ET (0)	TMC (0)		Rsvd (0)	LP (0)	
	3	Parameter Length (4h)								
	4	(MSB)		Parameter Value						(LSB)
	5									
	6									
	7									

### Drive Counters Page Fields:

**Note** The Current or Default Cumulative, or Current or Default Threshold values are returned according to the PC field of LOG SENSE.

The DU (Disable Update), DS (Disable Save), TSD (Target Save Disable), ET (Enable Threshold Comparison), TMC (Threshold Met Criteria) and LP (List Parameter) fields are collectively referred to as the *control byte*.

**Parameter** The supported Parameter Codes in this page are as follows:

Code	Param. Code	Description	Length (bytes)	Default Cumulative	Default Threshold
	0001h	Loads	4	0	FF FF FF FFh
	0002h	Write Drive Errors	4	0	FF FF FF FFh
	0003h	Read Drive Errors	4	0	FF FF FF FFh

**Loads:** This is the same as the Tape Log load count.

**Write Drive Errors:** The number of groups which could not be written successfully because of drive or mechanism failure.

**Read Drive Errors:** The number of groups which could not be read successfully because of drive or mechanism failure.

## Mechanism Counters Log Page

The Mechanism Counters Log Page has the following format. All fields in this log are initialized to zero in the following circumstances:

- At power-on
- After a SCSI reset
- After a tape load
- Following configuration for Error-Rate Testing
- After a user sends an request for initialization through LOG SELECT

		Bit								
		7	6	5	4	3	2	1	0	
Byte	0	Page Code (38h)								
	1	Reserved (0)								
	2	(MSB)		Page Length (A0h)						(LSB)
	3									
—Mechanism Counter Descriptors—										
	0	(MSB)		Parameter Code						(LSB)
	1									
	2	DU (0)	DS (1)	TSD (0)	ET (0)	TMC (0)		Rsvd (0)	LP (0)	
	3	Parameter Length (4h)								
	4	(MSB)		Parameter Value						(LSB)
	7									

### Mechanism Counters Page Fields:

**Note** The Current or Default Cumulative, or Current or Default Threshold values are returned according to the PC field of LOG SENSE.

The DU, DS, TSD, ET, TMC and LP fields are collectively referred to as the *control byte*.

**Parameter Code** The supported Parameter Codes in this page are as follows:

Param. Code	Description
0001h	Faulty 12V
0003h	Mode Sensor Fault
0004h	Tension Fault
0005h	Bad Diameter
0006h	Capstan Stalled
0007h	Drum Lock Lost
0008h	Drum Stalled
0009h	Drum Dref Lost
000Ah	Drum PG Lost
000Bh	Supply Reel Stuck Unthreading
000Ch	Supply Reel Stuck Capstan Mode
000Dh	Capstan Clean Slip
000Eh	Take-up Reel Stuck Capstan Mode
000Fh	RLS Stuck RL Mode
0010h	RAM Test Fail
0011h	ROM Check Fail
0012h	Supply Reel Brake Fail
0013h	Take-up Reel Brake Fail
0014h	Signal Processor Test

## Data Compression Transfer Log Page

This page allows the host to monitor the performance of DCLZ compression. The PC field in the LOG SENSE command determines whether current cumulative, default cumulative, or current/default threshold values are returned. The page has the following format:

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Page Code (39h)							
	1	Reserved (0)							
	2	(MSB)	Page Length (58h)						(LSB)
—Parameter Information—									
Byte	0	(MSB)	Parameter Code (n)						(LSB)
	1								
	2	DU (0)	DS (1)	TSD (0)	ETC (0)	TMC (00)	Rsvd (0)	LP (0)	
	3	Parameter Length (4h)							
	4	(MSB)	Parameter Value						(LSB)
7									

### Data Compression Log Page Fields:

Parameter Codes Supported Parameter Codes in this page and their meanings are as follows:

Parameter Codes	Code	Description
	0001h	Number of entities written
	0002h	Number of entities read
	0003h	Number of records written
	0004h	Number of records read
	0005h	Kilobytes to data compression
	0006h	Kilobytes from data compression
	0007h	Kilobytes to tape
	0008h	Kilobytes from tape
	0009h	Logical entity size
	000Ah	Physical entity size
	000Bh	Uncompressed entities

- Entities Written/Read** A record of the total number of complete entities written to or read from the tape since the last power-on or Clear Log operation.
- Records Written/Read** A record of the total number of records (both compressed and uncompressed) written to or read from the tape since the last power-on or Clear Log operation.
- Kilobytes to/from Data Compression** A record of the total number of kilobytes written to or read from the drive since the last power-on or Clear Log operation.
- Kilobytes to/from Tape** A record of the total number of kilobytes written to or read from the tape since the last power-on or Clear Log operation.
- Logical Entity Size** The logical size of the last entity written to or read from the tape:  
*size = number of records in entity × record size in bytes*
- Physical Entity Size** The physical size of the last entity written to or read from the tape:  
*size = entity header length + uncompressed data length*
- Uncompressed Entities** A record of the total number of times non-DCLZ entities have been encountered on the tape during read operations since the last power-on or Clear Log operation.

From these figures, the host can calculate the average compression ratio achieved by the drive since the last power-on or Clear Log operation as follows:

Commands  
LOG SENSE

$$\text{average compression ratio} = \frac{\text{kilobytes to data compression}}{\text{kilobytes to tape}}$$

Similarly, the host can calculate instantaneous compression ratio achieved by the drive for a particular read or write command:

$$\text{instantaneous compression ratio} = \frac{\text{logical entity size}}{\text{physical entity size}}$$

## Data Compression Trace Log Page

The Data Compression Trace log consists of four parameters, each representing an event which has occurred concerning the data compression code. The first entry is the oldest event, and the last the most recent.

Four parameters are always available even if the drive has just been powered up and the log is not filled. The oldest events are overwritten as new device events occur, so that logging never stops.

The PC field of the LOG SENSE command has no effect on the data returned in this log.

The log page has the following format:

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Page Code (3Ah)							
	1	Reserved (0)							
	2	(MSB)		Page Length (800h)					
	3							(LSB)	
—Data Compression Event Traces—									
Byte	0	(MSB)		Parameter Code (n)					
	1							(LSB)	
	2	DU (0)	DS (1)	TSD (0)	ETC (0)	TMC (00)	Rsvd (0)	LP (0)	
	3	Parameter Length (04h)							
	4	(MSB)		Trace Code					
	5							(LSB)	
	7								

## Channel Trace Log Page

The Channel Trace log consists of  $n$  parameters, each representing an event which has occurred concerning the channel code. The first entry is the oldest event, and the last the most recent.

$n$  parameters are always available even if the drive has just been powered up and the log is not filled. The oldest events are overwritten as new channel events occur, so that logging never stops.

The PC field of the LOG SENSE command has no effect on the data returned in this log.

The log page has the following format:

		Bit									
		7	6	5	4	3	2	1	0		
Byte	0	Page Code (3Bh)									
	1	Reserved (0)									
	2	(MSB)		Page Length							
	3									(LSB)	
—Channel Event Traces—											
Byte	0	(MSB)		Parameter Code ( $n$ )							
	1									(LSB)	
	2	DU (0)	DS (1)	TSD (0)	ET (0)	TMC (00)		Rsvd (0)	LP (0)		
	3	Parameter Length (02h)									
	4	(MSB)		Trace Code							
	5									(LSB)	

## MODE SELECT

15h

### Description:

MODE SELECT allows the host to modify certain parameters in the Mode Parameter pages in order to configure the drive. These parameters cannot be saved, and after a power-up or reset, the default configuration will always be set.

MODE SENSE allows the drive to report which configuration parameters it supports and what its current configuration is.

The following Mode Parameter pages are supported:

Code	Page
------	------

02	Disconnect-Reconnect page
0F	Data Compression Characteristics Mode page
10	Device Configuration page
11	Medium Partitions page, for defining and altering the number and size of partitions

Implementing MODE SELECT and MODE SENSE requires a certain amount of handshaking between the host and the drive. Before configuring the drive, the host should issue a MODE SENSE. The drive can then return its current configuration and indicate what parameters are configurable. The host interprets this information and may then issue MODE SELECT to set the drive to the host's preferred configuration.

The drive's behavior when it receives a MODE SELECT command is as follows:

- 1 Any unwritten data in the buffer is flushed to tape.
- 2 The Mode Select Parameter list is transferred from the host to the drive during the Data-Out phase. The list contains a header followed by zero or more Mode Parameter pages.
- 3 The HP C1533A sets internal parameters to the values in the pages. It checks values for validity as it goes along, so if the command terminates because of an error, earlier parameter values will already have been set to their new values.

---

**Note** In the MODE SELECT and MODE SENSE descriptions much of the information overlaps and the descriptions should be read in conjunction.

---

### Pre-Execution Checks:

Illegal Field    Flag Link    Bad LUN    Reservation    Deferred Error    Unit Attention

Only complete parts of a parameter list can be sent. It is illegal to send a partial mode parameter header, mode block descriptor, or mode page.

If the Medium Partitions page is sent, Parameter List, Media Access, Media Write, and Diagnostic Status checks are also performed. Note that the tape must be at BOP for this page to be sent.

If the Device Configuration page is sent, the Media Access check is performed.

Command Descriptor Block:

		Bit							
		7	6	5	4	3	2	1	0
0		Operation Code (15h)							
1		Logical Unit Number (0)			PF	Reserved (0)			SP (0)
2	Byte	Reserved (0)							
3		Reserved (0)							
4		Parameter List Length							
5		Reserved (0)						Flag	Link

CDB Fields:

<b>PF</b>	<b>Page Format</b>	<p>0 The Mode Select parameter data is not SCSI-2 mode page compatible. Only the parameter header and block descriptor may be sent. As a result, the HP C1533A will report ILLEGAL REQUEST if the Parameter List Length is greater than 12 bytes.</p> <p>1 The data sent by the host after the MODE SELECT header and block descriptors complies with the definition of pages in the SCSI-2 specification.</p>
<b>SP</b>	<b>Save Pages</b>	<p>0 This bit is not supported and must be zero.</p>
<b>Parameter List Length</b>	<p>0 No data is transferred.</p> <p>&gt;0 The length in bytes of the MODE SELECT parameter list to be transferred. The list length must be less than 4×number of pages.</p>	

Completion:

After successful completion of MODE SELECT, UNIT ATTENTION is posted to all initiators other than the initiator of the command. Additional sense is 2A01h (mode parameters changed).

Mode Select Specific Status:

Event	Status	Key	Additional Sense
Mode parameter rounded to match drive sensitivity	CHECK CONDITION	RECOVERED ERROR	3700h (rounded parameter)
Incorrect parameter list	CHECK CONDITION	ILLEGAL REQUEST	2000h (parameter list length error)
Parameter list too long	CHECK CONDITION	ILLEGAL REQUEST	1A00h (illegal field in Command Descriptor Block)

If the flush of write-behind data fails, this is reported as a deferred error.

---

## Mode Parameter Pages

### Mode Page Representation

Certain conventions are used in the following Mode Parameter pages in order to describe the nature of the parameters. Most parameters are given by name, followed by a number in brackets. The brackets have the following meanings:

**Square brackets [ ]** Square brackets indicate that the parameter may be modified. The number inside the brackets is the default value for the field—in other words, the power-up or reset value.

MODE SELECT may modify this value by sending the page with a new value in the field.

MODE SENSE (current values) will return the current value of the parameter.

MODE SENSE (default values) will return the value in brackets [ ].

MODE SENSE (changeable values) will return a value of all ones.

**Round brackets ( )** Round brackets indicate that the parameter is fixed. The number inside the brackets is the fixed value for the field.

MODE SELECT must set the parameter to this value, otherwise CHECK CONDITION will be reported with a sense key of ILLEGAL REQUEST.

MODE SENSE (current values) will return the fixed value.

MODE SENSE (default values) will return the fixed value.

MODE SENSE (changeable values) will return a value of all zeros.

**Curly brackets { }** Curly brackets indicate that the field is fixed. The number inside them will always be returned irrespective of whether current, default or changeable values are being returned.

### Parameter Rounding

The following parameters in the mode pages may be rounded by the HP C1533A:

- Disconnect Time Limit, to the nearest 100 microseconds
- Bus Inactivity Limit, to the nearest 100 microseconds
- Write Delay Time Limit, to the nearest 100 milliseconds

If this happens, CHECK CONDITION status is returned with a sense key of RECOVERED ERROR and an additional sense code of 3700h (rounded parameter), but the command is considered to have completed successfully.



## Mode Data Format

Mode data consists of a 4-byte header, optionally followed by block descriptor and Mode Parameter pages.

### Mode Parameter Header:

The Mode Parameter header has the following format:

		<i>Bit</i>							
		7	6	5	4	3	2	1	0
Byte	0	Mode Data Length							
	1	Medium Type (00h)							
	2	WP	Buffered Mode [001]				Speed (0h)		
	3	Block Descriptor Length							

### Header Fields:

<b>Mode Data Length</b>	<b>MODE SELECT:</b>	Mode Data Length field must be zero.
	<b>MODE SENSE:</b>	Mode Data Length field is set to the length of the data available. Note that the actual value returned may be truncated to the Allocation Length for the command.
<b>WP</b>	<b>Write Protect. Ignored for MODE SELECT.</b>	
	0	<b>MODE SENSE:</b> Tape is write-enabled
	1	<b>MODE SENSE:</b> Tape is write-protected
<b>Buffered Mode</b>	0	The drive will not report GOOD status on write commands until all blocks and filemarks have been written successfully.
	1	The HP C1533A will report GOOD status on write commands as soon as all the data has been transferred to the data buffer. Data for multiple commands from the different initiators may be buffered before being written to tape.
	2	The drive will report GOOD status on write commands as soon as both the following conditions are satisfied: <ul style="list-style-type: none"> <li>■ All data specified in the write command has been transferred to the data buffer.</li> <li>■ All buffered data from different initiators has been successfully written to tape.</li> </ul> Note that data for multiple commands from the same initiator may be buffered before being written to the medium.
<b>Block Descriptor Length</b>	Only the values 0 and 8 are valid.	
	0	No block descriptor is being transferred.
	8	The Mode Parameter header must be followed by the 8-byte block descriptor described in the next section.

## Mode Block Descriptor

The format of the Mode Parameter block descriptor is as follows:

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Density Code [13h]							
	1	(MSB)		Number of Blocks (0)					
	3							(LSB)	
	4	Reserved (0)							
	5	(MSB)		Block Length (0)					
	7							(LSB)	

### Mode Parameter Block Descriptor Fields:

<b>Density Code</b>	This must be 00h, 13h, 24h or 7Fh. For MODE SENSE 24h is returned for a 120m tape, 13h for 60m or 90m tapes. For MODE SELECT, the values have the following meanings: 00h Default (use this if you do not know the density) 13h HP/Sony DDS-1 and DDS-DC Formats—use this if the tape is 60m or 90m 24h HP/Sony DDS-2 Format—use this if the tape is 120m 7Fh No change from previous density (NO-OP)
<b>Block Length</b>	This indicates the size of fixed blocks. See the description of the Fixed fields in the READ and WRITE commands. It is also used when deciding whether to report an illegal length block on reads (see the SILI error in the Status section of the READ command).

## Disconnect-Reconnect Page

The drive supports the Disconnect-Reconnect Page which has the following format:

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Page Code (02h)							
	1	Page Length (0Eh)							
	2	Buffer Full Ratio (0)							
	3	Buffer Empty Ratio (0)							
	4	(MSB)		Bus Inactivity Limit [39h = 5.7 ms]					
	5							(LSB)	
	6	(MSB)		Disconnect Time Limit (0)					
	7							(LSB)	
	8	(MSB)		Connect Time Limit (0)					
	9							(LSB)	
	10	(MSB)		Maximum Burst Size [80h = 64 Kb]					
	11							(LSB)	
	12	Reserved (0)						DTDC (0)	
	13	Reserved (0)							
	14	Reserved (0)							
	15	Reserved (0)							

### Disconnect-Reconnect Page Fields:

Buffer Full Ratio, Buffer Empty Ratio, and DTDC (Data Transfer Disconnect Control) are not supported by the drive and must be cleared to zero.

<b>Bus Inactivity Limit</b>	<p>&gt;0 The maximum time in 100 microsecond increments for which the HP C1533A will assert the BSY signal without a REQ/ACK handshake. Since the HP C1533A's timing capabilities have a resolution of about 1 millisecond, parameter rounding may occur.</p> <p>Since the firmware overhead from the last ACK of a command phase to the first REQ of a data phase is typically 5 ms, setting the bus inactivity limit to &lt;5 ms will cause the firmware to disconnect following the command phase (typically &lt;700 <math>\mu</math>s from the last ACK of the command phase). This applies to a configurable subset of commands only, by default, read, write and spacing commands.</p> <p>0 The time limit becomes infinite. In practice, this means the HP C1533A will not disconnect between receipt of a command and the first data phase of that command. Note that a command with multiple data phases (for example, to match maximum burst size) will still have disconnects between phases.</p>
<b>Disconnect Time Limit</b>	The minimum time in 100 microsecond increments that the drive will wait after releasing the SCSI bus before attempting reselection. Since the HP C1533A's timing capabilities have a resolution of about 1 millisecond, parameter rounding may occur.
<b>Connect Time Limit</b>	The drive cannot guarantee to disconnect within a time limit.
<b>Maximum Burst Size</b>	<p>&gt;0 The maximum amount of data the drive will transfer during a data phase before disconnecting, expressed in units of 512 bytes.</p> <p>0 There is no limit to the amount of data that will be transferred.</p>

**Note** If the host does not grant disconnect privilege in its IDENTIFY message these parameters become meaningless. The HP C1533A will stay connected on the bus for the duration of an operation or until the host sends a mid-operation IDENTIFY granting disconnect privilege.

## Data Compression Characteristics Page

The Data Compression Characteristics Mode Page is defined as follows:

		<i>Bit</i>							
		7	6	5	4	3	2	1	0
<i>Byte</i>	0	Reserved (0)			Page Code (0Fh)				
	1	Page Length (0Eh)							
	2	DCE [1]	DCC (1)	Reserved (0)					
	3	DDE [1]	RED [0]		Reserved (0)				
	4	(MSB)	Compression Algorithm [20h]						(LSB)
	7								
	8	(MSB)	Decompression Algorithm [0]						(LSB)
	11								
	12	Reserved (0)							
	15								

### Data Compression Characteristics Mode Page Fields:

<b>DCE</b>	<i>Data Compression Enable</i>
<b>MODE SELECT</b>	This field allows the host to enable or disable data compression.
0	The drive will disable data compression, and any subsequent data sent to the drive by the host will be written to tape uncompressed.
1	The drive will compress any subsequent data sent to it by the host before writing it to tape. The Algorithm used to compress the data is that specified in the Compression Algorithm field.
<b>MODE SENSE</b>	This field allows the host to determine whether compression is enabled or disabled.
0	Compression is disabled
1	Compression is enabled
<b>DCC</b>	<i>Data Compression Capable</i>
<b>MODE SENSE</b>	This field allows the host to determine whether the drive supports data compression. This field is fixed.
0	Compression not supported
1	Compression is supported
<b>DOE</b>	<i>Data Decompression Enable</i>
<b>MODE SELECT</b>	This field allows the host to enable or disable decompression.
0	Decompression is disabled. Any compressed data encountered will be returned to the host as a single variable length record.
1	Decompression is enabled.
<b>MODE SENSE</b>	This field allows the host to determine whether the drive will attempt to decompress any entities it encounters on tape.
0	The drive will not attempt decompression
1	The drive will attempt compression

**RED** *Report Exception on Decompression*

This field allows the host to specify at which boundaries between different data types a CHECK CONDITION is to be reported to the host. The field only affects a READ command. In the following table, the data types are U=uncompressed, D=DCLZ compressed, non-D=non-DCLZ compressed. The different boundaries will generate CHECK CONDITION with the specified RED values.

Before	After	RED Values	Sense Key
U	non-D	0, 1, 2	MEDIUM SENSE
U	D	2	RECOVERED ERROR
D	U	2	NO SENSE
D	non-D	0, 1, 2	MEDIUM ERROR
non-D	U	1, 2	NO SENSE
non-D	D	1, 2	RECOVERED ERROR
non-D1	non-D2	0, 1, 2	MEDIUM ERROR

In each case, additional sense is 70nnh (decompression exception short algorithm ID of nnh). The tape is positioned on the EOP side of the boundary.

Decompression Algorithm field = that for the new data type  
Command-Specific Information field in the sense data = the number of data blocks contained within the new data.

<b>Compression Algorithm</b>		This indicates which compression algorithm the drive will use to process data from the host when the DCE bit is set to one. If the host selects an algorithm which the drive does not support, CHECK CONDITION is returned, with a sense key of ILLEGAL REQUEST. A value of zero indicates that no compression algorithm is currently selected. The drive only supports the DCLZ algorithm (ID 20h).
	<b>MODE SENSE</b>	The field contains the registered algorithm identifier for the currently selected compression algorithm. Note that this field will never be 1.
<b>Decompression Algorithm</b>	<b>MODE SELECT</b>	This field has no meaning and is ignored by the drive.
	<b>MODE SENSE</b>	This field allows the host to determine the type of compression algorithm used to process data sent most recently to the host in response to a READ command. <ul style="list-style-type: none"> <li>0 The last data item returned was an uncompressed record.</li> <li>n The last data item returned was an entity compressed using algorithm n. Note that this is valid, even if the data was decompressed by the drive.</li> </ul>

## Device Configuration Page

The drive supports the Device Configuration Page, which has the following format:

		Bit									
		7	6	5	4	3	2	1	0		
Byte	0	Page Code {10h}									
	1	Page Length {0Eh}									
	2	Rsvd (0)	CAP	CAF	RAW [0]	C3ECC[1]	N-Group {000}				
	3	Active Partition									
	4	Write Buffer Full Ratio (0)									
	5	Read Buffer Empty Ratio (0)									
	6	(MSB)		Write Delay Time {0032h}							
	7									(LSB)	
	8	DBR (0)	BIS (1)	RSmk [1]	AVC (0)	SOCF (00)		RBO (0)	REW [0]		
	9	Gap Size (0)									
	10	EOD Defined {000}			EEG (1)	SEW (0)	Reserved (0)				
	11	(MSB)		Buffer Size at Early Warning (0)							
	13									(LSB)	
	14	Select Data Compression Algorithm (0)									
	15	Reserved (0)									

### Pre-Execution Checks:

Before the page is sent in MODE SELECT data, the following checks are made:

Parameter List                      Media Access (if the CAP bit is set)

### Device Configuration Page Fields:

The DBR (Data Buffer Recovery), AVC (Automatic Velocity Control), SOCF (Stop On Consecutive Filemarks), RBO (Recover Buffer Order), Gap Size, Buffer Size at Early Warning and Select Data Compression Algorithm fields are not supported by the drive and must be cleared to zero.

<b>CAP</b>	<i>Change Active Partition</i>
	<b>MODE SELECT</b> If the CAP bit is set to 1, the logical partition is to be changed to that specified by the Active Partition field. If the new partition is the same as the current partition
	<b>MODE SENSE</b> The CAP bit must be zero.
<b>CAF</b>	<i>Change Active Format</i>
	<b>MODE SELECT</b> If the CAF bit is set to 1, the active format is to be changed to that given by the RAW, C3 ECC and N-Group fields.
	<b>MODE SENSE</b> The CAF bit must be zero.
<b>RAW</b>	<p>0 Read-After-Write is enabled. The drive rewrites frames which have not been read back correctly through RAW.</p> <p>1 Read-After-Write disabled. The results of RAW are ignored. The drive proceeds to write the next frame.</p> <p>To modify this bit using MODE SELECT, the CAF bit must be set.</p>
<b>C3 ECC</b>	<p>1 C3 ECC is enabled during writes, so that an ECC frame (frame 23) is written. This is the default configuration.</p> <p>0 C3 ECC is disabled, so frame 23 is not written.</p> <p>To modify this bit using MODE SELECT, the CAF bit must be set. The C3 ECC bit does not affect C3 ECC during reads. If a group has an ECC frame, C3 ECC will be performed on the group if necessary, irrespective of the value of this bit.</p>

<b>N-Group</b>	The number of times a group will be repeated on tape after the first writing. The default is zero repetitions. The more repetitions there are, the lower the data rate will be; this feature may be used to match the performance of the drive and the host. The range is 0 (000b) to 7 (111b). To modify this bit using MODE SELECT, the CAF bit must be set.
<b>Active Partition</b>	<p><b>MODE SELECT</b> Provided the CAP bit is set to 1, MODE SELECT will change the current partition to the partition indicated by this field. After the change, the tape will be positioned at BOP in the new partition. It follows that if you try to change the partition to the partition that is already active, the effect is simply of rewinding the tape. For a 1-partition tape, the only valid value is 00h. For a 2-partition tape, 00h and 01h are valid, and the default after load, power-up or reset is 00h. To format a tape as two partitions, use the Mode Select Medium Partitions page.</p> <p>Note that it is also possible to change the current active partition with the LOCATE command.</p> <p><b>MODE SENSE</b> This field indicates the current active partition.</p>
<b>Write Delay Time</b>	<p>0 The drive will never flush buffered data to tape as a result of a timeout.</p> <p>&gt;0 The maximum time, in 100 millisecond units, that the drive will wait with a partially full buffer before flushing unwritten data to tape.</p>
<b>DBR</b>	<i>Data Buffer Recovery.</i> This is always 0, indicating that DBR is not supported. This bit is ignored on MODE SELECT.
<b>BIS</b>	<i>Block Identifiers Supported.</i> This is always 1 to indicate that the drive has recorded information about the logical block ID relative to the partition. The bit is ignored on MODE SELECT.
<b>RSmk</b>	<p><b>Report Setmarks</b></p> <p>0 Setmarks are not reported, and are ignored.</p> <p>1 Setmarks are reported, that is, the drive will stop when it meets setmarks in certain circumstances. See the READ and SPACE commands for details. This is the default.</p>
<b>REW</b>	<i>Report Early-Warning End-of-Media.</i> For writes, the drive reports early-warning (EW) at a distance of 500 mm before EOP/M, irrespective of the state of the REW bit. For reads and space operations, early-warning will only be reported if this bit is set. See the READ and SPACE commands for details. The default is that the REW bit is cleared to zero.
<b>EOD Defined</b>	0 This field must be zero so that the drive will use its default EOD definition to detect and generate EOD.
<b>EEG</b>	1 <i>Enable EOD Generation.</i> This bit must be set to 1 to indicate that the drive will generate an EOD before any change of direction following a write-type operation. The bit is ignored on a MODE SELECT command.
<b>SEW</b>	0 <i>Synchronize at Early Warning.</i> The drive does not synchronize at the early warning point. However, the default firmware configuration permits a Mode Select command to set this bit without causing an error or affecting drive operation. This bit will always be 0 in the Mode Sense data.

## Medium Partitions Parameter Page

The drive supports the Medium Partitions Parameter Page, with the following format:

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Page Code {11h}							
	1	Page Length {08h} or {6h}							
	2	Maximum Additional Partitions {01h}							
	3	Additional Partitions Defined							
	4	FDP (0)	SDP (0)	IDP (1)	PSUM (10)		Reserved (0)		
	5	Medium Format Recognition {03h}							
	6	Reserved (0)							
	7	Reserved (0)							
	8	(MSB)	Size of Partition 1						(LSB)
9									

### Pre-Execution Checks:

Before the page is sent in MODE SELECT data, the following checks are made:

Parameter List                      Media Access                      Medium Write                      Diagnostic Status

### Medium Partitions Parameter Page Fields:

The FDP (Fixed Data Partitions) and SDP (Select Data Partitions) fields are not supported and must be cleared to zero.

<b>Page Length</b>	<b>MODE SENSE</b>	This will always return a page length of 8.
	<b>MODE SELECT</b>	This may return a page length of 6, providing the "Additional Partitions Defined" field is 0.
<b>Max. Additional Partitions</b>	01h	This field is only valid on MODE SENSE where it is used to report the maximum number of additional partitions that are supported by the drive, in this case, 1.
<b>Additional Partitions Defined</b>	0 and 1 are the only valid values. For MODE SENSE, the value is only valid if GOOD status is returned in response to a TEST UNIT READY command. If the page length is 6, this field must be 0, otherwise CHECK CONDITION is reported.	<ul style="list-style-type: none"> <li>0 For MODE SENSE this means the tape is one-partition. For MODE SELECT, it means the tape is to be formatted as one partition.</li> <li>1 For MODE SENSE this means the tape is two-partition. For MODE SELECT, it means the tape is to be formatted as two partitions.</li> </ul>
<b>IDP</b>	<i>Initiator Defined Partitions</i>	
	0	The remainder of the page is ignored. No action is taken.
	1	The drive formats the tape into the number and size of partitions given in the Additional Partitions Defined, PSUM and Partition Size fields. This is the default. See the Glossary for an explanation of partitions. Note that if a blank tape is loaded and not formatted before the first write command is received, the drive automatically formats the tape to have a single partition. It is necessary to send the Medium Partitions Parameter page only when creating a 2-Partition tape, or when reformatting a single partition tape to two partitions.
<b>PSUM</b>	<i>Partition Size Unit of Measure.</i> This field defines the units for the Partition Size value. For the HP C1533A, the units are megabytes, which is indicated by 10b. As a result, the drive returns 10b for MODE SENSE and will only accept this value for a MODE SELECT.	
<b>Medium Format Recognition</b>	03h	This must be set to 03h to indicate that the drive is capable of format and partition recognition.



---

<b>Size of Partition 1</b>	<b>If Page Length is 6, this field is not sent.</b>
<b>MODE SENSE</b>	<b>When the Additional Partitions Defined field is set to 1 (so that the tape is two-partition), the Size of Partition 1 field gives the size of partition 1 in megabytes. This is only valid when GOOD status is returned in response to TEST UNIT READY.</b>
<b>MODE SELECT</b>	<b>Size of Partition 1 indicates the size in megabytes for the drive to format Partition 1 in a 2-Partition tape. It must be 0 if the number of additional partitions is 0.</b>

# MODE SENSE

1Ah

## Description:

MODE SENSE allows the drive to return its current configuration and report which configuration parameters can be changed through MODE SELECT.

See MODE SELECT for full descriptions of the Mode Parameters.

**Note** if a MODE SENSE command is issued while a LOAD command with immediate reporting on is executing or near to completion, the values of some parameters in some mode pages which depend on the results of the load may be undefined. For example, the Medium Partitions page will only become valid after the load has completed.

Similarly, the WP (Write Protect) field in the mode parameter header will be unstable unless a cartridge is present and medium removal is prevented.

## Pre-Execution Checks:

Illegal Field    Flag Link    Bad LUN    Reservation    Deferred Error    Unit Attention

## Command Descriptor Block:

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

## CDB Fields:

<b>DBD</b>	<i>Disable Block Descriptors</i>		
	0	Allows the HP C1533A to return the Mode Select block descriptor.	
	1	Prevents the HP C1533A from returning the Mode Select block descriptor.	
<b>PC</b>	<i>Page Control</i> —Indicates the type of page parameter values to be returned to the host, as shown in the following table:		
	7	6	
	7	6	Description
	0	0	<i>Report Current Values</i> : this is the normal default situation where the drive returns its current configuration to the host. The drive sets any unsupported Page fields to zero.
	0	1	<i>Report Changeable Values</i> : these are any values which a host can alter through the MODE SELECT command. Any field that the drive allows to be changed is set to all ones, otherwise the field is cleared to zero.
	1	0	<i>Report Default Values</i> : these are the default values on power-on, reset, or bus-device reset.
	1	1	Invalid. A CHECK CONDITION will be generated.

---

**Page Code** This allows the host to select any specific page, or all the pages supported by the drive. The following page codes are supported:

Page Code	Description
00h	Only the 4-byte header and one 8-byte block descriptor are returned. This ensures compatibility with SCSI-1 as well as SCSI-2.
02h	Disconnect-Reconnect page returned
0F	Data Compression Characteristics Mode page returned
10h	Device Configuration Parameters page returned
11h	Medium Partition Parameters page returned
3Fh	All supported pages returned

The Additional Page Length field of each page returned by the drive indicates the number of bytes supported for that page. Fields not supported by the drive are set to zero.

---

**Allocation Length** Specifies the number of bytes that the host has allocated for returned MODE SENSE data. An Allocation Length of zero means that the drive will return no data. Any other value indicates the maximum number of bytes that can be transferred. The drive finishes the DATA IN phase when Allocation Length bytes have been transferred or when the entire mode page has been transferred to the host, whichever is less. See MODE SELECT for the format of the data returned by MODE SENSE.

# PREVENT/ALLOW MEDIUM REMOVAL

1Eh

## Description:

PREVENT/ALLOW MEDIUM REMOVAL tells the drive to enable or disable the removal of the cartridge.

## Pre-Execution Checks:

Illegal Field    Flag Link    Bad LUN    Reservation    Deferred Error    Unit Attention

## Command Descriptor Block:

		Bit							
		7	6	5	4	3	2	1	0
0		Operation Code (1Eh)							
1		Logical Unit Number (0)				Reserved (0)			
2	Byte	Reserved (0)							
3		Reserved (0)							
4		Reserved (0)							Prevent
5		Reserved (0)						Flag	Link

## CDB Fields:

The fields have the following meanings:

- |         |   |   |
|---------|---|---|
| Prevent | 0 | The drive ejects the tape following an UNLOAD command. The Unload button is also re-enabled.  |
|         | 1 | The drive will not eject the tape following an UNLOAD command. It will unthread the tape, but leaves the cartridge in the mechanism. The Unload button is disabled. |

It is recommended that the host sends a PREVENT MEDIUM REMOVAL command before it issues a LOAD command, so that the tape will not be unloaded immediately when the Unload or Eject button is pressed on the front panel, but only after the load is complete.

Removal prevention is specific to a host; only the host that imposed the prevention can clear it.

By default, medium removal is enabled. The drive resorts to this default state following a power-on, SCSI reset or bus device reset.

# READ

08h

## Description:

READ transfers zero or more data blocks to the host starting at the current position.

## Pre-Execution Checks:

Illegal Field	Fixed Bit	Flag Link	Bad LUN	Reservation
Deferred Error	Unit Attention	Media Access	Diagnostic Status	

## Command Descriptor Block:

		Bit								
		7	6	5	4	3	2	1	0	
0		Operation Code (08h)								
1		Logical Unit Number (0)				Reserved (0)			SIL	Fixed
2	(MSB)	Transfer Length								
3										
4										(LSB)
5		Reserved (0)						Flag	Link	

## CDB Fields:

SIL	Suppress Incorrect Length Indicator
0	The read operation is terminated when the length of a block (on the tape) differs from the Transfer Length. The number of bytes transferred will be the smaller of the transfer length requested and the length of blocks on the tape. The logical tape position will be the EOM side of the illegal sized block.
1	The HP C1533A will not report CHECK CONDITION status if the only error is that the Transfer Length is not equal to the actual block length recorded on the media. This applies for variable block mode reads only. <i>Exception:</i> If the block size in the mode block descriptor is not zero but the actual block length does not equal the requested block length, ILI will be reported in the sense data.
Fixed	0 The Transfer Length field specifies the length of the transfer in bytes. 1 The Transfer Length field specifies the length of the transfer in blocks. The size of each block (in bytes) is specified by the current block length specified in the Mode Parameter block descriptor.
Transfer Length	0 No data is transferred. This is not considered an error and the current logical position will be unchanged. >0 The amount of data to be transferred, in bytes or blocks as specified by the Fixed field.

The SIL and Fixed fields must not both be set, otherwise CHECK CONDITION is reported with a sense key of ILLEGAL REQUEST.

Read Specific Status:

Event	Status	Key	Notes
SILI error	CHECK COND'N	NO SENSE	ILI and Valid bits are set. The information bytes will be set as follows: <i>Variable Block Mode:</i> The difference between the requested Transfer Length and the actual block size. If the block size is bigger than the Transfer Length, this will be negative (twos complement). <i>Fixed Block Mode:</i> The difference between the required number of blocks and the number of whole good blocks transferred. In other words, the illegal length block is not counted.
Filemark encountered during a read	CHECK COND'N	NO SENSE	Mark bit is set. Logical position will be on the EOM side of the filemark. Additional sense is 0001h (filemark detected). Residue information is as described below.
Setmark encountered during a read	CHECK COND'N	NO SENSE	Mark bit is set. Report Setmarks must have been configured by setting the RSmk bit in the Mode Select Device Configuration page to 1. Additional sense is set to 0003h (setmark detected). The logical position is on the EOM side of the setmark. Residue information is as described below.
EOD encountered during a read	CHECK COND'N	BLANK CHECK	Additional sense is set to 0005h (EOD detected). Residue information is as described below.
EOP/M encountered during a read	CHECK COND'N	MEDIUM ERROR	EOM bit is set. Additional sense is set to 0002h (EOP/M detected). Residue information is as described below.
Early warning (EW) EOM detected during a read	CHECK COND'N	NO SENSE	Additional sense set to 0002h (EOP/M detected). Early warning end of medium (EW) information is reported only if the read operation was successful and the REW bit in the Mode Select Device Configuration page is set to 1.
Failure to read data through media error or non-fatal drive error	CHECK COND'N	H/W ERROR	Additional sense is set to 3B00h (sequential positioning error). Residue information is as described below.
Incompatible data format	CHECK COND'N	MEDIA ERROR	Additional sense set to 3002h (can't read media).
Blank media encountered during a read	CHECK COND'N	BLANK CHECK	Additional sense is set to 1403h (EOD not found). Residue information is as shown below.

Residue Information:

The Valid bit will be set. The information bytes will be as follows:

**Variable Block Mode** The difference (in bytes) between the requested Transfer Length and the length of a block on the tape.

**Fixed Block Mode** The difference (in blocks) between the requested transfer size and the actual number of blocks transferred. A partially transferred block is not counted, so for example, if 3.5 out of 10 blocks have been transferred the residue will be 7.

If the last block is an Illegal Length block, it is not counted among those transferred. For example, if 4 Legal Length blocks and 1 Illegal Length block have been transferred out of a total of 10, the residue will be 6.

# READ BLOCK LIMITS

**05h**

**Description:**

**READ BLOCK LIMITS** tells the drive to return its limits for the block length it can support. The **READ BLOCK LIMITS** data shown below is sent during the **DATA IN** phase of the command. The command does not reflect the currently selected block size, only the available limits. **MODE SENSE** returns the current block size.

**Pre-Execution Checks:**

Illegal Field    Flag Link    Bad LUN    Reservation    Deferred Error    Unit Attention

**Command Descriptor Block:**

		<i>Bit</i>							
		7	6	5	4	3	2	1	0
<i>Byte</i>	0	Operation Code (05h)							
	1	Logical Unit Number (0)				Reserved (0)			
	2	Reserved (0)							
	3	Reserved (0)							
	4	Reserved (0)							
	5	Reserved (0)						Flag	Link

**Read Block Limits Data:**

		<i>Bit</i>									
		7	6	5	4	3	2	1	0		
<i>Byte</i>	0	00h									
	1	(MSB)		Maximum Block Length							
	2	unsigned (FF FF FFh)									
	3									(LSB)	
	4	(MSB)		Minimum Block Length (0001h)							
	5									(LSB)	

# READ BUFFER

3Ch

## Description:

READ BUFFER returns a segment of the drive's data buffer, mini buffer, processor memory, or a buffer descriptor to the host during the Data-In phase. It is used in conjunction with WRITE BUFFER as a diagnostic function for testing the 1 megabyte data buffer and the SCSI bus integrity of the drive.

**Note** This command cannot be used to recover buffered data. If the drive receives a READ BUFFER without having had a prior WRITE BUFFER command, only the 4-byte header will be returned, because the drive writes any buffered data to tape before accepting READ BUFFER or WRITE BUFFER for the first time. However, if READ BUFFER is sent *immediately* after WRITE BUFFER, all or part of the data just written can be recovered.

## Pre-Execution Checks:

Illegal Field      Flag Link      Bad LUN      Reservation      Deferred Error      Unit Attention

## Command Descriptor Block:

		BK							
		7	6	5	4	3	2	1	0
0		Operation Code (3Ch)							
1		Logical Unit Number (0)			Reserved (0)			Mode	
2		Buffer ID							
3	(MSB)	Buffer Offset							
5									
6	(MSB)	Allocation Length							
8									
9		Reserved (0)						Flag	Link

## CDB Fields:

<b>Mode</b>	0	<i>Read header and data.</i> The HP C1533A returns a 4-byte buffer header followed by buffer data starting from address 0 of the data buffer (if Buffer ID = 0), or the mini buffer (if Buffer ID = 2). The header is identical to the buffer descriptor (see Buffer Descriptor on the next page). The total length of header plus data is Allocation Length
	1	<i>Read processor memory.</i> The HP C1533A returns Allocation Length bytes of data from an address in processor memory specified by Buffer Offset.
	2	<i>Read data buffer or mini buffer.</i> The HP C1533A returns Allocation Length bytes of data from the data buffer (if Buffer ID is 0) or the mini buffer (if Buffer ID is 1). The start address of the data in the buffer is specified by Buffer Offset.
	3	<i>Return buffer descriptor.</i> The HP C1533A returns a descriptor for the buffer identified by the Buffer ID field. For buffer 0, the data buffer, the Buffer Descriptor (see the next page) is returned with Buffer Capacity of 08 00 00h. For buffer 2, the Buffer Descriptor is returned with Buffer Capacity 02 00 00h. For all other buffer IDs, the Buffer Capacity is returned as zero.
See the Valid Field Values section below for constraints upon the value of Mode imposed by the Buffer ID field.		
<b>Buffer ID</b>	0	Read the data buffer
	1	Read the processor memory
	2	Read the mini buffer



<b>Buffer Offset</b>	When Mode = 2, the Buffer Offset specifies the address in the data buffer or mini buffer from which to start reading, relative to the start of the buffer. When Mode = 1, the Buffer Offset specifies the absolute address in processor memory from which to start reading.
<b>Allocation Length</b>	When Mode = 1, the Allocation Length specifies the number of bytes of data to be transferred from processor memory. When Mode = 2, it is the number of bytes to be transferred from the data buffer or mini buffer. If the combined length and offset exceeds the size of the buffer, the drive wraps around from the buffer end and returns the remainder from the start of the buffer. When Mode = 0, a 4-byte header is returned, followed by (Allocation Length - 4 bytes of data starting at address 0 in the data buffer. If Allocation Length is less than 4, nothing is returned. When Mode = 3, Allocation Length is ignored; the 4-byte descriptor alone is returned.

**Valid Field Values:**

The following table shows the values of the Buffer ID, Buffer Offset and Allocation Length fields for each valid value of Mode:

Mode	Buffer ID	Buffer Offset	Allocation Length
If Mode = 0 then:	0	0	≤ 8 00 00h + 4
	2	0	≤ 2 00 00h + 4
If Mode = 1 then:	1	Within valid controller processor address range	So that the window of processor memory starting at Buffer Offset is valid
If Mode = 2 then:	0	< 8 00 00h	≤ 8 00 00h
	2	< 2 00 00h	≤ 2 00 00h
If Mode = 3 then:	0 or 1	0	Ignored

**Returned Data:**

The data returned by READ BUFFER depends on the setting of the Mode field. Buffer Descriptor A below shows the 4-byte header returned as part of the data for the Combined Header and Data mode (Mode 0). This is identical to the data returned for Return Buffer Descriptor mode (Mode 3) for the data buffer (Buffer ID = 0).

Note that READ BUFFER cannot be used to read buffered data. Data read from the buffer is undefined unless the command is immediately preceded by a WRITE BUFFER command.

**Buffer Descriptor**

		Bk							
		7	6	5	4	3	2	1	0
Byte	0	00h							
	1	(MSB)							
	2	Buffer Capacity							
	3								(LSB)

- For Buffer ID 0, a Buffer Capacity of 8 00 00h is returned.
- For Buffer ID 2, a Buffer Capacity of 2 00 00h is returned.
- For all other Buffer IDs, a Buffer Capacity of zero is returned.

**Read Buffer Specific Status:**

If READ BUFFER fails, a sense key of HARDWARE ERROR is returned with additional sense code of 4001h (diagnostic failure on component 1). This could be because the hardware has failed, or it could happen when a previously unwritten part of the data buffer is read. This is because the parity bits related to the bytes being read are undefined at power-on, so they cause spurious parity errors.

# READ POSITION

34h

**Description:**

The READ POSITION command returns the current logical position of the drive to the host. The position is given in two parts:

- The current partition.
- The offset from the start of the partition in terms either of blocks (ignoring marks), or the count of all blocks and marks. The BT bit selects which. Block 0 is located at the start of the partition.

**Pre-Execution Checks:**

Illegal Field Deferred Error	Flag Link Unit Attention	Bad LUN Media Access	Reservation Diagnostic Status
---------------------------------	-----------------------------	-------------------------	----------------------------------

**Command Descriptor Block:**

		<i>Bit</i>							
		7	6	5	4	3	2	1	0
<b>0</b>		Operation Code (34h)							
<b>1</b>		Logical Unit Number (0)				Reserved (0)			BT
<b>2</b>	<i>Byte</i>	Reserved (0)							
:		:							
<b>8</b>		Reserved (0)							
<b>9</b>		Reserved (0)						Flag	Link

**CDB Fields:**

<b>BT</b>	<i>Block address Type</i>
0	The First Block Location field is the total number of blocks, filemarks and setmarks from the beginning of the partition.
1	The value in the First Block Location field is to be the number of blocks from the beginning of the partition to the current position.

Returned Data:

The format of the 20 bytes of data returned for READ POSITION is as follows:

		Bit								
		7	6	5	4	3	2	1	0	
0	BOP	EOP		Reserved (0)			BPU	Reserved (0)		
1	Partition Number									
2	Reserved (0)									
3	Reserved (0)									
4	(MSB)	First Block Location								
7									(LSB)	
8	(MSB)	Last Block Location								
11		(same as First Block Location)							(LSB)	
12	Reserved (0)									
13	(MSB)	Number of Blocks in Buffer (0)								
15									(LSB)	
16	(MSB)	Number of Bytes in Buffer (0)								
19									(LSB)	

Returned Data Fields:

<b>BOP</b>	0	The current logical position is not the beginning of the partition.
	1	The current logical position is the beginning of the partition.
<b>EOP</b>	0	The current logical position is not between early warning and the end of the current partition
	1	The current logical position is between early warning and the end of the current partition
<i>Note that the EOP and BOP bits are mutually exclusive.</i>		
<b>BPU</b>	<i>Block Position Unknown</i>	
	0	First and Last Block Location fields are valid.
	1	First and Last Block Location fields are undefined. This will only occur after a "hard" read error which has caused a sense key of MEDIUM ERROR with additional sense of 3B00h (sequential positioning error).
<b>Partition Number</b>	The value of the current partition.	
	0	Valid for 1- or 2-Partition tapes.
	1	Valid for 2-Partition tapes only.
<b>First Block Location</b>	The First Block Location specifies either the number of blocks, or the total number of blocks plus marks from the beginning of the partition to the current position. The BT bit determines which measure is used. Counting starts at zero, so BOP is block location 0.	
<b>Last Block Location</b>	This is always set to the value of the First Block Location.	

# RECEIVE DIAGNOSTICS RESULTS

1Ch

## Description:

RECEIVE DIAGNOSTIC RESULTS tells the drive to send analysis data to the host after completion of a SEND DIAGNOSTIC command.

Data is specific to the last SEND DIAGNOSTIC command received by the drive, whatever host sent it. If no SEND DIAGNOSTIC has been received, the data refers to the power-on self-test.

## Pre-Execution Checks:

Illegal Field    Flag Link    Bad LUN    Reservation    Deferred Error    Unit Attention

## Command Descriptor Block:

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Operation Code (1Ch)							
	1	Logical Unit Number (0)				Reserved (0)			
	2	Reserved (0)							
	3	(MSB)		Allocation Length					
	4							(LSB)	
	5	Reserved (0)						Flag	Link

## CDB Fields:

<b>Allocation Length</b>	0	No data is to be returned.
	>0	The number of bytes which the host has allocated for returned diagnostic data. The HP C1533A returns Allocation Length bytes or what is available, whichever is less. The amount of data available is always 64 bytes.

## Results:

See the Diagnostics section in this manual for information on the returned data.

# RELEASE UNIT

17h

## Description:

**RELEASE UNIT** enables the host to release any reservation it may have on the drive. Reserving the drive is a way of ensuring that the host has exclusive access to the drive.

## Pre-Execution Checks:

Illegal Field      Flag Link      Bad LUN      Deferred Error      Unit Attention

## Command Descriptor Block:

		Bit							
		7	6	5	4	3	2	1	0
0		Operation Code (17h)							
1		Logical Unit Number (0)			3rd Pty	Third-Party Device ID			Rsvd (0)
2		Reserved (0)							
3		Reserved (0)							
4		Reserved (0)							
5		Reserved (0)						Flag	Link

## CDB Fields:

3rdPty	0	The host wishes to release a reservation it has made for itself.
	1	The host wishes to release a reservation it has made for a third-party host.
Third-Party Device ID		If the 3rdPty bit is set to 1, the 3rd Party Device ID field specifies the ID of the device for which the HP C1533A has been reserved by the host. If the 3rd Pty bit is clear, this field is ignored.

## Release Unit Specific Status:

Status is returned as follows:

This Host	The HP C1533A	Another Host or 3rd Party	3rd Pty bit	3rd Party ID	Status	Action Taken
	not reserved				GOOD	None
	reserved by...	another host			GOOD	None
reserved drive for itself	reserved		Clear		GOOD	Cleared
reserved drive for itself	reserved		Set		GOOD	None
reserved drive for 3rd Pty host	reserved for...	3rd Party host	Clear or	<> ID of 3rd Party Host	GOOD	None
reserved drive for 3rd Pty host	reserved for...	3rd Party host	Set	= ID of 3rd Party Host	GOOD	Cleared

# REQUEST BLOCK ADDRESS

02h

## Description:

REQUEST BLOCK ADDRESS returns a count representing the current logical block address to the host. A block is of fixed or variable length depending on the Fixed bit in the WRITE command. Blocks are numbered sequentially, starting with 1 at BOM. Each block, filemark, or setmark adds one to the count from BOM.

## Pre-Execution Checks:

Illegal Field Deferred Error	Flag Link Unit Attention	Bad LUN Media Access	Reservation Diagnostic Status
---------------------------------	-----------------------------	-------------------------	----------------------------------

## Command Descriptor Block:

		<i>Bit</i>							
		7	6	5	4	3	2	1	0
<i>Byte</i>	0	Operation Code (02h)							
	1	Logical Unit Number (0)				Reserved (0)			
	2	Reserved (0)							
	3	Reserved (0)							
	4	Allocation Length							
	5	Reserved (0)						Flag	Link

## CDB Fields:

<b>Allocation Length</b>	0	The default transfer length of 3 bytes is returned. The MSB is returned first.
	>0	This number of bytes is returned to the host, up to a maximum of 3 bytes.

# REQUEST SENSE

03h

## Description:

The REQUEST SENSE command transfers sense data held within the drive to the host during a data-in phase. The data is valid in the following circumstances:

- After reporting CHECK CONDITION status to the host.
- After a command has terminated with an unexpected BUS FREE.
- Immediately after the following commands which cause tape motion:  
     READ      WRITE      WRITE FILEMARK      REWIND

**Note** The drive maintains valid positional information for normal read and write commands only (READ, WRITE, SPACE, WRITE FILEMARK, REWIND). If abnormal commands (such as WRITE BUFFER) are received this positional sense information is lost.

The drive clears sense data for the host following execution of the REQUEST SENSE command for that host, unless the data is positional sense data.

## Pre-Execution Checks:

Illegal Field      Flag Link      Bad LUN

## Command Descriptor Block:

		<i>Bit</i>							
		7	6	5	4	3	2	1	0
<i>Byte</i>	0	Operation Code (03h)							
	1	Logical Unit Number (0)				Reserved (0)			
	2	Reserved (0)							
	3	Reserved (0)							
	4	Allocation Length							
	5	Reserved (0)							Flag

## CDB Fields:

<b>Allocation Length</b>	<b>0</b>	No data is transferred to the host.
	<b>&gt;0</b>	The maximum amount of sense data in bytes that should be transferred to the host. The drive returns up to Allocation Length bytes of data; any extra sense data is lost.

## Request Sense Data

The following diagram shows the format of the data returned for the REQUEST SENSE command. The HP C1533A always has 19 bytes of sense data available to return. Note that separate copies of sense data are maintained for the 7 possible hosts that could share the SCSI bus with the drive.

The returned sense information has the following format:

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Valid		Error Code (70h or 71h)					
	1	Segment Number (0)							
	2	Mark	EOM	ILJ	Rsvd (0)	Sense Key			
	3	(MSB) Information Bytes							
	6	(LSB)							
	7	Additional Sense Length (normally 08h, up to 2Ch if COPY ABORT sense)							
	8	(MSB) Command Specific Information Bytes (0)							
	11	(LSB)							
	12	Additional Sense Code							
	13	Additional Sense Code Qualifier							
	14	FRU Code							
	15	SKSV	C/D	Reserved (0)		BPV	Bit Pointer		
	16	(MSB) Field Pointer/Drive Error Code							
	17	(LSB)							
	18-51	Reserved (0)							

Returned Sense Data Fields:

Valid	1	Indicates that the information bytes contain valid information as defined in the SCSI specification.
Error Code	70h	Indicates that the error is <i>current</i> , that is, it is associated with the command for which CHECK CONDITION status has been reported.
	71h	Indicates that the error is <i>deferred</i> . A deferred error occurs when there is a failure for an operation which has already been terminated with a GOOD status, or when failure occurs in "cleanup" activity following an operation which was terminated by BUS RELEASE. The command for which CHECK CONDITION status was reported is therefore unlikely to be the cause of the deferred error. It simply gives the HP C1533A the opportunity to report CHECK CONDITION status for an error which already exists.
Segment Number	This field is cleared to zero.	
Mark	1	The Mark bit is set to 1 when the logical position is at a filemark or setmark. See the READ and SPACE commands. The Additional Sense Code and Qualifier fields will indicate which sort of mark it is. The Mark bit is only set when the sense key is NO SENSE. Setmarks will only be reported if the RSmk bit in the MODE SELECT Device Configuration Parameters page is set.
EOM	1	The EOM bit indicates the tape is at BOP or EOP/M. Three different sense keys may be set: NO SENSE is set when the tape is at BOP with additional sense 0004h, or in the Early-Warning region preceding EOP with additional sense 0002h. BOP can be reported after a SPACE command encounters BOP, or after commands such as LOAD and REWIND. EOP/M is reported when a WRITE or WRITE FILEMARK command leaves the tape positioned in the early warning (EW) region near the end of the partition. It is also reported when a READ or SPACE command actually hits EOP/M. See the READ, SPACE, WRITE and WRITE FILEMARKS commands. See the Glossary for a description of EW regions, BOP and EOP/M. VOLUME OVERFLOW is reported when physical EOP is encountered during a write operation. BLANK CHECK is reported when physical EOP is encountered during a read or space operation.
ILJ	1	Indicates that the requested block length did not match the actual block length. Only READ and VERIFY can cause this bit to be set.



<b>Sense Key</b>	See "Sense Keys and Additional Sense Codes" later in this section for details of sense keys which can be returned.
<b>Information Bytes</b>	The Information Bytes field is only valid if the Valid bit is set. It contains residue information following the failure of either a READ, WRITE, WRITE FILEMARKS or SPACE command. The field can take a negative value, expressed in two's complement notation. See the description of the command that failed for the contents of the field.
<b>Additional Sense Length</b>	0Bh Specifies the number of additional sense bytes to follow and is always set to 0Bh. It is never truncated to reflect actual transfer length.
<b>Command Specific Information Bytes</b>	0 The normal value for this field. >0 If the entity is compressed with an algorithm other than DCLZ, this field contains the number of records in the entity.
<b>Additional Sense Code and Qualifier</b>	These two bytes provide additional information about the cause of the CHECK CONDITION, or the current tape position when the REQUEST SENSE is unsolicited. See "Sense Keys and Additional Sense Codes" later in this section.
<b>FRU Code</b>	This will be set to a value to indicate a specific part which has failed. It may take one of the following values: 0 No failing unit 1 Controller PCA failure 2 Mechanism failure
<b>SKSV</b>	0 The Sense Key Specific Bytes (bytes 15 through 17) are invalid. The C/D, BPV and Bit pointer fields will be zero. The top byte of the Field Pointer will be zero, and the bottom byte will contain a product-specific error code. 1 The Sense Key Specific Bytes are valid. This will only happen when the bytes assume the role of Field Pointer Bytes as defined by ANSI. This occurs when an Illegal Field check detects an error in a command descriptor block or a Parameter List check detects an error in a command parameter list.
<b>C/D</b>	<i>Command Data</i> . Only valid if the SKSV flag is set. 0 The Field Pointer information applies to the parameter list for the command. 1 The Field Pointer information applies to the command descriptor block.
<b>BPV</b>	<i>Bit Pointer Valid</i> 0 The Bit Pointer field is invalid. 1 The Bit Pointer field is valid. If the SKSV bit is set then the BPV bit must be set as well.
<b>Bit Pointer</b>	Identifies the bit position of the field in error, whether it is a command descriptor or a command parameter list field. Only valid when the SKSV flag is set.
<b>Field Pointer / Drive Error Code</b>	If SKSV is set, this field identifies in which byte of the command descriptor or parameter list an error was detected. Note that the HP C1533A scans from byte 0 bit 7 of a command or parameter list through to byte <i>n</i> bit 0. The field and bit pointers are set to point to the most significant bit of the field that is in error. If SKSV is not set, the top byte of this field will be zero. The bottom byte contains the drive error code associated with the failure of the previous command, or zero in the case of no failure and no appropriate error code. <i>Note:</i> Drive error codes are provided for information only; they are intended to assist in product integration and fault diagnosis. The codes should <i>not</i> be relied on by "generic" host drivers, since these are specific to the product rather than the particular revision of the product.

---

## Sense Data Management

The HP C1533A maintains three sets of sense data for every host on the bus. For a single host the following sense data is maintained:

- Current sense
- Unit Attention sense
- Deferred Error sense

Unless otherwise stated, all the following descriptions apply to the sense data for a single host whose command is being executed.

### Current Sense

*Current sense* is the data which is returned in response to a REQUEST SENSE command.

Current sense is modified or updated in the following circumstances:

- It is cleared in response to the arrival of any command other than REQUEST SENSE or INQUIRY. It is cleared following execution of a REQUEST SENSE command. It is cleared on arrival of an unsolicited REQUEST SENSE command (that is, when no CHECK CONDITION status was reported for the previous command). This ensures the positional sense is updated to reflect current position before returning sense data.
- If the current command fails, sense is set according to the failure and CHECK CONDITION is reported to the host which sent the command.
- If the previous command terminated with an unexpected BUS FREE, current sense is set according to the reason for the bus release.
- If the command fails its pre-execution check for Unit Attention, Unit Attention sense is copied to the Current sense, and CHECK CONDITION is reported to the host which sent the command.
- If the command fails its pre-execution check for Deferred Error, Deferred Error sense is copied to the Current sense, and CHECK CONDITION is reported to the host which sent the command.

### Clearing Current Sense:

Clearing the Current sense involves setting sense fields as follows:

Valid	0	Error Code	70h	Segment Number	0
ILI	0	Sense Key	0	Information Bytes	0
Command Specific Information	0	FRU Code	0	SKSV	0

### Unsolicited Sense:

Normal sense data is available to a host through the REQUEST SENSE command only when the host's previous command was reported with CHECK CONDITION status. A REQUEST SENSE command is *unsolicited* when there was no preceding CHECK CONDITION, and positional sense information may be set.

If Unsolicited Positional Sense (a compile option) is disabled, the only sense that can be reported is BOP or Early-Warning EOP.

This unsolicited sense indicates the current logical and physical position of the media. The sense key will be set to NO SENSE. The Mark and EOM bits will be set depending on whether the tape is at a mark or at end of media. The additional sense will be set to indicate the current position (at a mark or at EOM).

Unsolicited sense regards marks as of a higher priority than early warning. As a result, if it meets a mark in the Early-Warning region, the drive might set a CHECK CONDITION with additional sense of "end of medium/partition detected", but subsequent requests for sense would result in the additional sense code and qualifier changing to "filemark/setmark encountered". The EOM flag would still be set.

---

**Note** If the HP C1533A is performing activities related to a command with immediate report such as REWIND, unsolicited positional sense will give the position *before* the rewind. When the rewind operation genuinely completes, unsolicited positional sense will indicate BOM.

---

*Example 1:*

The drive tries to space "behind" BOP. CHECK CONDITION status is reported with a sense of "BOP encountered". If a subsequent unsolicited REQUEST SENSE command is sent, the sense remains as "BOP encountered" irrespective of the configuration of unsolicited sense reporting.

*Example 2:*

The drive meets a filemark while reading in the Early-Warning region. It reports CHECK CONDITION status, with both Mark and Com flags set. A subsequent unsolicited REQUEST SENSE command returns the same sense data, providing that unsolicited positional sense reporting is enabled. Otherwise, the subsequent REQUEST SENSE will return sense data indicating "early warning end of media".

## Unit Attention Sense

Unit Attention sense is set when one of the following Unit Attention conditions occur:

- Power On, SCSI Reset, Bus Device Reset. Additional sense will be set to 2900h (power-on, reset or bus clear reset occurred).
- Tape Loaded (possible media change). Additional sense will be set to 2800h (not ready to transition). See LOAD/UNLOAD command Specific Status.
- Mode Sense Parameters Changed. Additional sense will be set to 2A01h (mode parameters changed). See MODE SENSE command Specific Status.
- Microcode Downloaded. Additional sense will be set to 3F01h (microcode downloaded). See WRITE BUFFER command.

Unit Attention sense will persist for a host until the host sends a command which has Unit Attention as part of its pre-execution checks (that is, any command other than INQUIRY and REQUEST SENSE). When the Unit Attention pre-execution check fails, the Unit Attention sense is moved to the Current sense and CHECK CONDITION status is reported to the host. The host is then expected to issue a REQUEST SENSE command to recover the new Current sense data.

It is possible for multiple Unit Attention conditions to arise (for example, on power-on followed by tape load) so that the host does not read one condition before the next occurs. In this case a prioritizing scheme is used. If a Unit Attention condition exists and a new one occurs, then the Unit Attention sense will be overwritten by the new sense only if the new sense has a higher priority. The order of priority is as follows:

Power-on, Reset	<i>highest priority</i>
Tape Loaded	
Mode Parameters Changed	
Microcode Downloaded	<i>lowest priority</i>

## Deferred Error Sense

Deferred Error sense is generated when an operation fails and CHECK CONDITION status cannot be reported to the host for one of the following reasons:

- The operation had immediate-report on (for example, REWIND with the Immed bit set, or LOAD).
- The operation was a write which was immediate-reported on after its data was successfully placed in the write-behind queue.
- The operation was abandoned without status (that is, a bus release occurred) after what the HP C1533A perceived was a catastrophic error or when an error occurred following the HP C1533A's responding to an ABORT message. Following the bus release, the drive performed "cleanup" activity which failed. The Deferred Error refers to the reason for this failure.

When any one of the above occurs, the sense generated will be Deferred Error sense. This sense persists until the host sends a command which has Deferred Error as part of its pre-execution checks (that is, any command other than INQUIRY and REQUEST SENSE). When the Deferred Error pre-execution check fails, the Deferred Error sense is moved to the Current sense and CHECK CONDITION status is reported to the host. The host is then expected to issue a REQUEST SENSE command to recover the new Current sense data.

### Write-Behind Deferred Errors:

If the Buffered Mode field of the MODE SELECT parameter header is not zero, data buffering is enabled. If an error occurs in flushing the write-behind data, Deferred Error sense is set for all those hosts which the HP C1533A thinks have write-behind data in the buffer.

For each host with this sense set, the following occurs when its next command is to be executed:

- If the command is WRITE or WRITE FILEMARKS then the Deferred Error sense is copied to Current sense and the error is changed to a Current error (that is, the Error Code field is set to 70h). The residue information will reflect the size of the write-behind queue plus the transfer size of the current command. CHECK CONDITION status is reported. When the host reads sense data with a REQUEST SENSE command it will see an error for the WRITE command with residue information indicating data lost for that write and previous writes.
- If the command is not a write command and the Deferred Error check is performed as part of the pre-execution checks then CHECK CONDITION status is reported. The Deferred Error sense is copied to Current sense. The Deferred Error sense is cleared. The host is expected to retrieve the sense by sending a REQUEST SENSE command.

### Bus Release Deferred Errors:

The HP C1533A will release the bus to terminate an operation when the host sends an ABORT message or when the HP C1533A fails to send status due to a catastrophic error (see MESSAGE OUT OVERFLOW, STATUS PHASE ABORT and ILLEGAL PARITY ERROR messages).

In either case, the operation (for example, READ) does not terminate immediately since the logical position will have to be established on the other side of the block that was about to be transferred, or was being transferred. If an error subsequently occurs while trying to complete the read, then the sense generated is a Deferred Error. It is therefore possible but highly unlikely to have Deferred Error sense flagged for a command without immediate report.

## Sense Keys and Additional Sense Codes

This section contains a list of all the sense keys and additional sense codes that the HP C1533A can return. Along with each code is a description of the cause of the sense.

### 0h NO SENSE

Indicates that there is no specific sense key information to be reported. This is the case for a successful command, or a command that has received CHECK CONDITION or COMMAND TERMINATED status because of the filemark, EOM, or ILI bits is set to one.

Code	Description
00 02h	<b>EOP/M detected</b> A WRITE or WRITE FILEMARK command completed successfully leaving the tape at the early warning point of the current partition.
00 03h	<b>Setmark detected</b> Unsolicited positional sense has been set to indicate "At a Setmark". See READ, SPACE, and WRITE FILEMARKS commands. The Mark bit in the sense data will always be set when this additional sense is returned. or The command failed because a setmark was detected during READ, SPACE, or VERIFY. This is only possible if setmark reporting has been enabled in the Mode Select Device Configuration page. The Mark bit in the sense data will always be set when this additional sense is returned.
00 04h	<b>BOP/M detected</b> A SPACE command encountered BOP unexpectedly. Note that unless unsolicited positional sense indicates BOP, a request to format the tape through MODE SELECT Medium Partitions will be rejected. or Unsolicited positional sense has been set to report "At BOP". This occurs after a successful REWIND command. It also occurs after clearing CHECK CONDITION status when a SPACE has met BOP, and after clearing UNIT ATTENTION status after a successful LOAD. Note: unless unsolicited positional sense indicates BOP, a request to format the tape using MODE SELECT Medium Partitions will be rejected.
00 05h	<b>End of data detected</b> Unsolicited positional sense has been set to indicate "At EOD". See SPACE, WRITE, and WRITE FILEMARKS commands.

### 1h RECOVERED ERROR

Indicates that the last command completed successfully with some recovery action performed by the target.

Code	Description
15 01h	<b>Mechanism positioning error</b> This is set whenever the drive has performed a mechanism retry in successfully completing a command (READ, WRITE, WRITE FILEMARKS, SPACE or ERASE).
17 00h	<b>Recovered data with no error correction applied</b> This is set when an unspecified "soft" error occurs while successfully completing a command (READ, WRITE, WRITE FILEMARKS, SPACE or ERASE).
17 01h	<b>Recovered data with retries</b> This is set if one or more frames has to be rewritten by the drive in order to complete a command successfully (WRITE, or WRITE FILEMARKS).
18 00h	<b>Recovered data with error correction applied</b> This is set when the drive uses C3 ECC to complete a command successfully (READ, or SPACE).

---

## 2h NOT READY

---

Indicates that the logical unit being addressed cannot be accessed.

Code	Description
04 00h	<b>Logical unit not ready, cause not reportable</b> This is set if a LOAD/UNLOAD command is received while an unload is occurring with immediate report on, or initiated through the front panel. Note that for non-load commands, the "unloading" status is reported by additional sense 04 02h.
04 01h	<b>Logical unit in process of becoming ready</b> This is set if a media access command is received while a load with immediate report on, or autoloading, is occurring.
04 02h	<b>Logical unit not ready, Initializing command required</b> This is set if a command is received and a tape is present but unloaded, or being unloaded. A LOAD command is required to access the tape.
30 02h	<b>Incompatible media</b> This is set when incompatible media is detected by the Medium Recognition System.
37 00h	<b>Rounded parameter</b> This status is generated during execution of a MODE SELECT command, where the HP C1533A has to round a mode parameter specified by the host to match the drive's sensitivity.
3A 00h	<b>Medium not present</b> This is set when a command which accesses the tape is received and no cartridge is present in the drive.
3E 00h	<b>Logical unit has not self-configured yet</b> This is set during power-on when it is not possible to send commands which access the media to the drive because mechanism tests are being executed. This additional sense changes on completion of the mechanism tests to either 3A00h, 0401h, or 0402h depending on whether a cartridge was present at power-on and whether autoloading is configured.

---

## 3h MEDIUM ERROR

---

Indicates that the command terminated with a non-recovered error condition which was probably caused by a flaw in the tape or an error in the recorded data.

Code	Description
00 02h	<b>EOP/M detected</b> A READ, SPACE, VERIFY, WRITE or WRITE FILEMARK command encountered EOP/M unexpectedly.
11 00h	<b>Unrecovered read error</b> This is set if a command fails because a tape read failed. The most likely cause is bad media, but it could result from a hardware error.
30 01h	<b>Cannot read media, unknown format</b> This is set if a READ or SPACE command encounters an unknown tape format.
30 03h	<b>Cleaning cartridge installed</b>
33 00h	<b>Tape length error</b> This is set if a tape is found which is too short to format with a partition of the size requested.
38 00h	<b>Sequential positioning error</b> The HP C1533A could not read a group off tape. There are two possibilities: The current command (such as READ, SPACE, or REWIND) failed to complete successfully. or The logical position has been lost. The tape has been positioned on the other side of the bad groups. It is up to the host to determine if it is worth continuing. Note that there is a danger of skipping EOD and running onto old data if the host continues. The host must decide whether the data is old or current.

3B 01h	<b>Tape position at BOM</b> This is set if a LOAD command or a LOCATE/MODE SENSE command to change partitions fails to read the System area of the tape.
50 00h	<b>Write append error</b> This is set if a WRITE or WRITE FILEMARKS command fails because the append point was unreadable
52 00h	<b>Cartridge fault</b> This is set if the drive thinks there is something wrong with the cartridge. Possible causes are snapped tape, or an invalid combination of identification holes.
53 00h	<b>Media load or eject failed</b> This is set if a load or eject fails.
53 01h	<b>Unload tape failure</b> This is set if an UNLOAD operation fails.
70 nnh	<b>Decompression exception short algorithm ID of nnh</b> This is set during a read command when the drive detects a change of recorded data type (record or entity) which it is not configured to return to the host. This additional sense code is also returned if the target block address for a SPACE or LOCATE command is within an entity that the drive cannot decompress.

#### 4h **HARDWARE ERROR**

Indicates that the target detected a non-recoverable hardware failure while performing the command or during self-test.

Code	Description
09 00h	<b>Track following error</b> This is set if a read fails due to a fatal positioning error. The error could be related to the tape or the drive.
0C 00h	<b>Write error</b> The HP C1533A has failed to write data or marks to tape. The error could be related to the tape or the device. Note that separate additional sense is reported in the case of an ERASE failure. Residue information will normally be supplied.
31 00h	<b>Format command failed</b> This is set if a formatting operation induced through a MODE SELECT command sending a Media Partitions page fails.
40 XXh	<b>Diagnostic failure on component XX</b> This is set when a self-test command detects an error or when a command is prohibited because of the failure of a diagnostic.
53 00h	<b>Media load or eject failed</b> This is set if a load or eject fails.

#### 5h **ILLEGAL REQUEST**

Indicates that there was an illegal parameter in the command descriptor block, or in additional parameters.

Code	Description
20 00h	<b>Invalid command operation code</b> This is set if the HP C1533A does not recognize the opcode of the command it has received.
24 00h	<b>Invalid field in Command Descriptor Block</b> This is set if the HP C1533A detects an invalid field during the pre-execution checks in a command it has received.
25 00h	<b>Logical unit not supported</b> This is set if an IDENTIFY message is received with a Logical Unit other than zero specified.
26 00h	<b>Invalid field in parameter list</b>

This is set if the HP C1533A detects an invalid field while scanning command parameters.

---

**53 02h Medium removal prevented**

This is set if an UNLOAD command is received while a tape is present but not loaded and medium removal has been prevented.

---

**6h UNIT ATTENTION**

---

Indicates that the tape may have been changed, or that the target has been reset.

---

**Code Description**

**28 00h Not ready to transition**

This is set following a tape load to indicate that the tape may have changed.

**29 00h Power-on, reset or bus device reset occurred**

This is set following a power-on, SCSI Reset signal or bus device reset message.

**2A 01h Mode parameters changed**

UNIT ATTENTION is set for all hosts following a MODE SELECT command, other than the host that issued the command.

**3F 01h Microcode downloaded**

This is set if a WRITE BUFFER command is used to download and execute code from the host. UNIT ATTENTION sense is posted to all hosts other than the current one

**53 00h Media load or eject failed**

This is set if a load or eject fails.

---

**7h DATA PROTECT**

---

Indicates that a write command was attempted on a write-protected block

---

**Code Description**

**27 00h Write-protected**

This is set if an operation which writes data, writes marks, or erases data, is requested for a write-protected cartridge.

---

**8h BLANK CHECK**

---

Indicates that the drive encountered blank tape or EOD while reading.

---

**Code Description**

**00 05h End of data detected**

A READ, SPACE, or VERIFY command encountered EOD unexpectedly.

**14 03h End of data not found**

This is set when the drive encounters blank tape while reading. The most likely cause is a corrupt format tape.

**3D 00h Invalid bits in Identify message**

This is set when the HP C1533A receives an illegal Identify message.

**43 00h Message error**

This is set when the number of parity errors or problems on inbound or outbound messages exceeds the limit. It is likely that a bus release will occur if excessive parity errors are detected because the HP C1533A will fail to complete the Status phase and COMMAND COMPLETE message. In this case, a Deferred Error is posted.

**46 00h Select or Reselect failure**

This is set when the drive fails to reselect a host to complete an operation. A Check Condition state will be assumed to exist despite the fact that CHECK CONDITION was never reported.

**49 00h Invalid message error**

This is set for a number of reasons related to the host and the HP C1533A not recognizing each others' messages.

---



<b>4A 00h</b>	<b>Command phase error</b> This is set if parity errors occur during an attempted Command phase.
<b>4B 00h</b>	<b>Data phase error</b> This is set if parity errors occur during the Data-In and Data-Out phases of an operation.
<b>4E 00h</b>	<b>Overlapped commands attempted</b> This is set when a host selects the HP C1533A when a command already exists for that host.

---

### Other

---

The following additional sense codes can occur with a variety of sense keys.

<b>Code</b>	<b>Description</b>
<b>00 00h</b>	<b>No additional sense information</b> This is returned when the HP C1533A has no additional sense information for the host.
<b>0A 00h</b>	<b>Media warning</b> This is returned when the drive detected an excessive number of RAW retries or C3 ECC errors. The condition can be caused by dirty heads, so the heads should be cleaned and the operation repeated. If the warning is repeated, it should be assumed that the cartridge is nearing the end of its useful life. The cartridge should be discarded, after any data on it has been copied.
<b>44 00h</b>	<b>Internal target failure</b> This code is used to report errors related to the firmware, when the HP C1533A encounters an "impossible" situation. This should never occur.
<b>47 00h</b>	<b>SCSI parity error</b> This is set when the HP C1533A detects a parity error in an unexpected SCSI state. It should never do this, so this error should be treated as a firmware error.
<b>48 00h</b>	<b>Initiator detected error message received</b> This is set if an INITIATOR DETECTED ERROR message is received and the previous phase (the phase in which ATN was asserted) was invalid.

# RESERVE UNIT

16h

## Description:

This command enables the host to reserve the HP C1533A. Reserving a device is a way of ensuring exclusive access to that device from a single host for the period of the reservation.

Once reserved, the HP C1533A will execute commands received from the host which reserved it or from a third-party host if the 3rd Pty option has been selected. Commands from other hosts, other than INQUIRY, REQUEST SENSE and RELEASE UNIT, will have RESERVATION CONFLICT reported for them. The INQUIRY and REQUEST SENSE commands are immune to the effects of a reservation and will continue to execute for all hosts. The RELEASE UNIT command will have GOOD status reported for other hosts, but will have no effect on the reservation.

If a third-party reservation is in effect, the host which reserved the drive may only issue INQUIRY, REQUEST SENSE, RESERVE and RELEASE commands.

If the HP C1533A is already reserved by another host, RESERVATION CONFLICT is returned.

If a host which has reserved the drive sends another RESERVE UNIT, the command is still effective so that the host can alter its reservation.

The reservation will stay in effect until:

- The reserving host sends another RESERVE UNIT command.
- The reserving host sends a RELEASE UNIT command, clearing the reservation.
- A power-on, SCSI Reset or Bus Device Reset occurs.

## Pre-Execution Checks:

Illegal Field      Flag Link      Bad LUN      Deferred Error      Unit Attention

## Command Descriptor Block:

		Bit								
		7	6	5	4	3	2	1	0	
0		Operation Code (16h)								
1		Logical Unit Number (0)			3rd Pty	Third-Party Device ID			Rsvd (0)	
2	Byte	Reserved (0)								
3		Reserved (0)								
4		Reserved (0)								
5		Reserved (0)							Flag	Link

## CDB Fields:

3rdPty	0	The host wishes to reserve the HP C1533A for its own use.
	1	The host wishes to reserve the HP C1533A for a third-party host.
Third-Party Device ID	If the 3rdPty bit is 0, the Third-Party Device ID field is ignored. If the 3rdPty bit is 1, the Third-Party Device ID field specifies the ID of the host that wishes to reserve the HP C1533A.	

# REWIND

01h

## Description:

REWIND tells the drive to position the tape at the beginning of the currently active partition. Before rewinding, the drive writes any write-behind data in the buffer to tape. Once a rewind is started, it will complete even if the SCSI command is aborted for some reason.

## Pre-Execution Checks:

Illegal Field	Flag Link	Bad LUN	Reservation
Deferred Error	Unit Attention	Media Access	Diagnostic Status

## Command Descriptor Block:

		<i>Bit</i>							
		7	6	5	4	3	2	1	0
<i>Byte</i>	0	Operation Code (01h)							
	1	Logical Unit Number (0)				Reserved (0)			Immed
	2	Reserved (0)							
	3	Reserved (0)							
	4	Reserved (0)							
	5	Reserved (0)						Flag	Link

## CDB Fields:

<b>Immed</b>	<b>0</b>	Status is returned after the rewind has completed.
	<b>1</b>	The drive first writes any unwritten buffered data to tape. It then returns GOOD status to the host before beginning the actual rewind operation.

## Rewind Specific Status:

If the rewind is successful, unsolicited positional sense will indicate that the tape is at BOP by the EOD bit being set and an additional sense code of 0004h (BOP).

# SEEK BLOCK

0Ch

## Description:

SEEK BLOCK is used to position the tape at a specified position in the current partition, either before or after the current position. When there is no exception condition during the execution of the command, the tape is positioned at the BOP side of the block specified in the Block Address field. The first block in a partition is numbered "1", so an attempt to seek block 0 will be considered an error.

## Pre-Execution Checks:

Illegal Field Deferred Error	Flag Link Unit Attention	Bad LUN Media Access	Reservation Diagnostic Status
---------------------------------	-----------------------------	-------------------------	----------------------------------

## Command Descriptor Block:

		Bit							
		7	6	5	4	3	2	1	0
0		Operation Code (0Ch)							
1		Logical Unit Number (0)				Reserved (0)			Immed
2	Byte	(MSB) Block Address							
3									
4									
5		Reserved (0)						Flag	Link

## CDB Fields:

<b>Block Address</b>	The block address to seek. Blocks are numbered sequentially, starting at one. All records, variable or fixed, filemarks, and setmarks add one to the block address.
----------------------	---

# SEND DIAGNOSTIC

1Dh

## Description:

SEND DIAGNOSTIC tells the drive to perform either its standard self-test, or diagnostics specified in the parameter list. The parameter list is transferred to the drive in a 8-byte data-out phases; each specifying a diagnostic to be performed by the drive. It is not an error to set Self-Test = 0 and to have a zero-length parameter list.

The results can be retrieved using the RECEIVE DIAGNOSTIC RESULTS command.

## Pre-Execution Checks:

Illegal Field    Flag Link    Bad LUN    Reservation    Deferred Error    Unit Attention

## Command Descriptor Block:

		Bit							
		7	6	5	4	3	2	1	0
0		Operation Code (1Dh)							
1		Logical Unit Number (0)		PF (0)	Rsvd (0)	Self-Test	Dev OffL(0)	Unit OffL	
2		Reserved (0)							
3	(MSB)	Parameter List Length							
4									
5		Reserved (0)						Flag	Link

## CDB Fields:

PF	Page Format
0	Parameters conform to SCSI-1 (all parameters are vendor-specific).
Self-Test	0 Perform the diagnostics specified in the parameter list. The Unit OffL bit must be set, otherwise CHECK CONDITION is reported.
	1 Perform the standard self-test. The Parameter List Length must be 0, otherwise CHECK CONDITION is reported.
Unit OffL	0 Indicates that the host is not prepared for the drive to perform tests which could affect logical position.
	1 Indicates that the host is prepared for the drive to perform tests which could affect logical position, modify the tape or modify cartridge status. This bit must be set if the Self-Test bit is 0.
Parameter List Length	This should be 0 or a multiple of 8. If the Self-Test bit is set, the Parameter List Length must be 0, otherwise CHECK CONDITION is reported. A zero length Parameter List is permitted when the Self-Test bit is clear.

SEND DIAGNOSTIC is documented more fully in the section covering Diagnostics.

## Send Diagnostic Specific Status:

If all the diagnostics pass, then GOOD status is reported.

Condition	Status	Sense Key	Additional Sense
A diagnostic fails on component <i>x</i>	CHECK CONDITION	HW ERROR	400xh

A diagnostic can only "fail" if its Continue of Sequence Error flag is clear. If this flag is set for a diagnostic in the parameter list, and the diagnostic fails to complete successfully, the drive will continue executing the remaining diagnostics in the parameter list. If the flag is clear, the drive will not execute any diagnostics remaining in the list.

# SPACE

11h

**Note** In this command description the word “mark” covers filemarks and setmarks. “Filemark” and “setmark” are only referred to explicitly when the spacing activity differs.

### Description:

SPACE provides a variety of positioning functions that are determined by Code and Count fields in the Command Descriptor Block. Both forward (towards EOM) and reverse (towards BOM) positioning are provided.

Any unwritten data in the buffer is flushed to tape before the space is started. Once the space has started, it will complete even if the SCSI operation is aborted.

### Pre-Execution Checks:

Illegal Field Deferred Error	Flag Link Unit Attention	Bad LUN Media Access	Reservation Diagnostic Status
---------------------------------	-----------------------------	-------------------------	----------------------------------

### Command Descriptor Block:

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Operation Code (11h)							
	1	Logical Unit Number (0)				Reserved (0)		Code	
	2	Count (n)							
	3								
	4								
	5								
5	Reserved (0)						Flag	Link	

### CDB Fields:

<b>Code</b>	<p>The Code field indicates what is to be spaced to (<i>n</i> is the value in the Count field):</p> <ul style="list-style-type: none"> <li>000 <b>Blocks</b>—Space over <i>n</i> blocks.</li> <li>001 <b>Filemarks</b>—Space over <i>n</i> filemarks.</li> <li>010 <b>Sequential Filemarks</b>—Space until a sequence of <i>n</i> consecutive filemarks is found. Logical position on reporting will be on the further side of the sequence of filemarks.</li> <li>011 <b>EOD</b>—Space to the end of data in the current active partition.</li> <li>100 <b>Setmarks</b>—Space over <i>n</i> setmarks. The RSmk (Report Setmark) bit of the MODE SELECT Device Configuration page must be set for this option.</li> <li>101 <b>Sequential Setmarks</b>—Space until a sequence of <i>n</i> consecutive setmarks is found. Logical position on reporting will be on the further side of the sequence of setmarks. The RSmk (Report Setmark) bit of the MODE SELECT Device Configuration page must be set for this option.</li> </ul>
<b>Count</b>	<p>When spacing over <i>blocks</i> or <i>marks</i> this field indicates how many blocks or marks should be crossed. Spacing is forward (towards EOM) unless Count is negative (twos complement) when the spacing is towards BOM.</p> <p>When spacing over <i>sequential marks</i>, this field indicates how many consecutive marks should be encountered. Spacing is forward (towards EOM) unless Count is negative (twos complement) when the spacing is towards BOM.</p> <p>When spacing to <i>EOD</i> the Count field is ignored.</p> <p>When not spacing to EOD, if Count is zero, the SPACE command is treated as a null operation. The logical position remains unchanged, and unwritten data is not flushed to tape.</p>

### Space Specific Status:

The HP C1533A implements the following priority scheme during spacing:

#### Lowest priority Blocks

- Filemarks** If the drive meets a filemark while spacing to a block CHECK CONDITION is reported. The Mark bit in the sense data is set, and the sense key is NO SENSE. Additional sense is 0001h (filemark encountered). The final position will be before or after the filemark depending on whether the space was reverse or forward.
- Setmarks** If the drive meets a setmark while spacing to a block or a filemark, and the RSmk bit has been set in the Active Format field of the Mode Select Device Configuration page, CHECK CONDITION is reported. The Mark bit in the sense data is set. The sense key is NO SENSE and additional sense is 0003h (setmark encountered). The final position will be before or after the setmark depending on whether the space was reverse or forward.
- EOD** If the drive meets EOD while spacing to a block or mark, CHECK CONDITION is reported. If it meets EOD after the EW mark, the EOM bit in the sense data is set. The sense key is BLANK CHECK and additional sense is 0005h (EOD encountered).

**Highest priority BOP, EOP/M** If the drive meets BOP or EOP/M while spacing, CHECK CONDITION is reported. The EOM bit in the sense data is set. The sense key is NO SENSE and additional sense is 0004h (BOP detected) or 0002h (EOP/M detected).

If "blank media" is encountered during a space (in other words, the drive attempted to space on an unformatted tape), the drive behaves as if EOD was at BOP. It will then respond with GOOD status to a request to space to EOD, and with CHECK CONDITION to any other forward spacing request. Sense data will be set as if EOD had been encountered.

Event	Status	Key	Notes
Early warning (EW) EOM detected	CHECK COND'N	NO SENSE	Additional sense set to 0002h (EOP/M detected). Early warning end-of-medium (EW) information is reported only if the space operation was successful and the REW bit in the Mode Select Device Configuration page is set to 1.
Incompatible data format	CHECK COND'N	MEDIUM ERROR	Additional sense will be set to 3001h (can't read media).
Fail to read data due to a tape error or non-fatal drive error	CHECK COND'N	H/W ERROR	Additional sense will be set to 3B00h (sequential positioning error).
Blank media encountered	CHECK COND'N	BLANK CHECK	Additional sense is set to 1403h (EOD not found). Residue information is as shown below.
Host requests a target block address which corresponds to a record which lies within an entity which the drive cannot decompress	CHECK COND'N	MEDIUM ERROR	Additional sense is set to 3B00h (sequential positioning error)

### Residue Information:

If SPACE fails when spacing to blocks or marks, the sense data Information Bytes will be set to the absolute value of the difference between the requested number of marks and the actual number of marks spaced over. The residue for a space towards BOP which terminates prematurely will always be a positive value. This value will only be accurate if the sense key is NO SENSE.

**Note** It is possible to configure Signed Residue reporting for space operations so that the Information Bytes will be negative (twos complement) if the space was towards BOP. This is *not* done by default.

# TEST UNIT READY

00h

## Description:

TEST UNIT READY checks if the drive is ready for commands which access the tape. It is not a request for a self-test. If the drive has a tape loaded, the command returns a GOOD status. Otherwise, CHECK CONDITION is reported and the sense key is NOT READY.

## Pre-Execution Checks:

Illegal Field  
Deferred Error

Flag Link  
Unit Attention

Bad LUN  
Media Access

Reservation

## Command Descriptor Block:

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Operation Code (00h)							
	1	Logical Unit Number (0)				Reserved (0)			
	2	Reserved (0)							
	3	Reserved (0)							
	4	Reserved (0)							
	5	Reserved (0)							Flag

## Test Unit Ready Specific Status:

If all the pre-execution checks pass, then GOOD status is reported.



# VERIFY

13h

## Description:

VERIFY verifies one or more blocks beginning with the next block on the tape. The verification is media verification only. Data is read from the tape as it would be for a read, but is immediately discarded. No data is transferred between the host and drive. VERIFY has no SILI bit. If an illegal length block is encountered, VERIFY behaves like READ with the SILI bit clear.

VERIFY will not complete until all data has been read. It does not support the ANSI "byte compare" or "immed" options.

## Pre-Execution Checks:

Illegal Field	Fixed Bit	Flag Link	Bad LUN	Reservation
Deferred Error	Unit Attention	Media Access	Diagnostic Status	

## Command Descriptor Block:

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Operation Code (13h)							
	1	Logical Unit Number (0)				Reserved (0)			Fixed
	2	(MSB)	Verification Length						
	3								
	4								(LSB)
	5	Reserved (0)						Flag	Link

## CDB Fields:

<b>Fixed</b>	0	The Verification Length field specifies the length of the data to be verified in bytes.
	1	The Verification field specifies the length of the data to be verified in blocks. The size of each block (in bytes) is specified by the current block length specified in the Mode Parameter block descriptor.
<b>Verification Length</b>	0	No data is verified. This is not considered an error and the current logical position will be unchanged.
	>0	The amount of data to be verified, in bytes or blocks as specified by the Fixed field.

Verify Specific Status:

Event	Status	Key	Notes
Filemark encountered during a read	CHECK COND'N	NO SENSE	Mark bit is set. Logical position will be on the EOM side of the filemark. Additional sense is 0001h (filemark detected). Residue information is as described below.
Setmark encountered during a read	CHECK COND'N	NO SENSE	Mark bit is set. Report Setmarks must have been configured by setting the RSmk bit in the Mode Select Device Configuration page to 1. Additional sense is set to 0003h (setmark detected). The logical position is on the EOM side of the setmark. Residue information is as described below.
EOD encountered during a read	CHECK COND'N	BLANK CHECK	Additional sense is set to 0005h (EOD detected). Residue information is as described below.
EOP/M encountered during a read	CHECK COND'N	MEDIUM ERROR	EOM bit is set. Additional sense is set to 0002h (EOP/M detected). Residue information is as described below.
Early warning (EW) EOM detected during a read	CHECK COND'N	NO SENSE	Additional sense set to 0002h (EOP/M detected). Early warning end of medium (EW) information is reported only if the read operation was successful and the REW bit in the Mode Select Device Configuration page is set to 1.
Failure to read data through media error or non-fatal drive error	CHECK COND'N	H/W ERROR	Additional sense is set to 3B00h (sequential positioning error). Residue information is as described below.
Incompatible data format	CHECK COND'N	MEDIA ERROR	Additional sense set to 3002h (can't read media).
Blank media encountered during a read	CHECK COND'N	BLANK CHECK	Additional sense is set to 1403h (EOD not found). Residue information is as shown below.

Residue Information:

The Valid bit will be set. The information bytes will be as follows:

- Variable Block Mode** The difference (in bytes) between the requested Verification Length and the actual transfer length.
- Fixed Block Mode** The difference (in blocks) between the requested size of the data to be verified and the actual number of blocks verified. A partially verified block is not counted, so for example, if 3.5 out of 10 blocks have been verified the residue will be 7.

If the last block is an Illegal Length block, it is not counted among those verified. For example, if 4 Legal Length blocks and 1 Illegal Length block have been verified out of a total of 10, the residue will be 6.

# WRITE

0Ah

## Description:

Zero or more blocks of data are transferred from the host to tape starting at the current logical position.

It is recommended that the Buffered Mode field of the MODE SELECT Parameter is set to either 1 or 2 so that Immediate Report is enabled. Failure to do this could result in a severe performance and capacity penalty. The HP C1533A will report GOOD status on a write command when all the data has been successfully transferred to the data buffer (but not necessarily to tape).

Data held in the buffer is flushed to tape in the following circumstances:

- A SCSI command is received which forces a flush:
 

ERASE	LOAD	LOG SELECT	MODE SELECT	READ
REWIND	SEND DIAGNOSTIC	SPACE	VERIFY	WRITE BUFFER
- Buffered mode is not enabled. If buffered mode has not been set to 1 or 2, the buffer is flushed before the command completes.
- The write delay time is exceeded. This is defined by the Write Delay field of the Device Configuration Mode Select page. If the drive is idle for longer than this period (no operations which access the tape have been performed), any data in the buffer is flushed to tape.
- A new host takes control of the drive. If buffered mode is set to 2, a WRITE or WRITE FILEMARKS command from the new host will cause all buffered data from the previous host to be flushed to tape.

## Pre-Execution Checks:

Illegal Field	Fixed Bit	Flag Link	Bad LUN	Reservation
Deferred Error	Unit Attention	Media Access	Media Write	Diagnostic Status

## Command Descriptor Block:

		Bit															
		7	6	5	4	3	2	1	0								
Byte	0	Operation Code (0Ah)															
	1	Logical Unit Number (0)				Reserved (0)				Fixed							
	2	Transfer Length															
	3									(MSB)							
	4																
	5	Reserved (0)						Flag	Link								

## CDB Fields:

<b>Fixed</b>	0	The Transfer Length field specifies the length of the transfer in bytes.
	1	The Transfer Length field specifies the length of the transfer in blocks. The size of each block (in bytes) is determined by the current Block Length given in the Mode Select Parameter block descriptor.
<b>Transfer Length</b>	0	No data is transferred. The current logical position will be unchanged.
	>0	The amount of data to be transferred, in bytes or blocks as determined by the Fixed field.

Write Specific Status:

Event	Status	Key	Notes
EOP/M encountered	CHECK COND'N	MEDIUM ERROR	Additional sense will be set to 0002h (EOP/M detected). The logical position will be at EOD. Subsequent unsolicited REQUEST SENSE commands however will give the tape position as EOP/M detected. It may still be possible to write a smaller quantity of information at this point (for example, filemarks, setmarks or a smaller block). This is because the logical position after failing to write a block is at the start of the unwritten block. Residue information is as below.
Early warning end-of-medium (EW) encountered	CHECK COND'N	NO SENSE	This is only reported if the write has been successful. Additional sense is set to 0002h (EOM/P detected). Residue information is valid and indicates zero blocks or bytes.

Residue Information:

Residue information depends on two variables:

- Fixed or Variable block mode.
- Immediate or Non-Immediate Report (Buffered mode).

Block Mode	Buffered Mode	Residue
Variable	0	Number of unwritten bytes
Fixed	0	Number of unwritten blocks
Variable	not 0	Total number of unwritten bytes and marks (including those buffered before the command was received). The size can be greater than the command operation size.
	specifically 1	Contains residues of all initiators with data in the buffer.
Fixed	not 0	Total number of unwritten blocks and marks (including those buffered before the command was received). The size can be greater than the command operation size.
	specifically 1	Contains residues of all initiators with data in the buffer.

Residue information is set in the sense data byte, with the Valid bit set and the information bytes set to the residue.

# WRITE BUFFER

3Bh

## Description:

WRITE BUFFER is used with READ BUFFER as a diagnostic for testing the 1 megabyte buffer and the SCSI bus integrity of the HP C1533A. The data is placed into the data buffer, the mini buffer or the HP C1533A's controller memory depending on the Mode and Buffer ID fields of the command. When data is to be placed in the data buffer or mini buffer, the contents of the buffer are flushed to tape first.

This command is intended as a diagnostic aid to be used in conjunction with READ BUFFER. It is not possible to modify data in the buffer after it has been written to tape.

## Pre-Execution Checks:

Illegal Field    Flag Link    Bad LUN    Reservation    Deferred Error    Unit Attention

## Command Descriptor Block:

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

## CDB Fields:

<b>Mode</b>	0	<i>Write header and data.</i> The HP C1533A transfers a 4-byte buffer header from the host which it ignores, followed by ( <i>Allocation Length</i> - 4) bytes of data which it will place in the buffer selected by the Buffer ID at address 0.
	1	<i>Write processor memory.</i> The HP C1533A transfers Allocation Length bytes of data from the host and writes them in processor memory specified by Buffer Offset.
	2	<i>Write data.</i> The HP C1533A transfers Allocation Length bytes of data from the host, placing them in the data buffer (if Buffer ID is 0) or mini buffer (if Buffer ID is 2). The start address of the data in the buffer is specified by Buffer Offset.
	4	<i>Download Microcode.</i> The HP C1533A transfers Allocation Length bytes of data from the host. If the transfer is successful, UNIT ATTENTION is posted to all initiators other than the current one, with additional sense set to 3F01H (microcode downloaded). The microcode is then executed.
See the Status section below for constraints upon the value of Mode imposed by the Buffer ID field.		
<b>Buffer ID</b>	Specifies which buffer to read.	
	0	The data buffer
	1	Processor memory
	2	Mini buffer
<b>Buffer Offset</b>	When Mode = 2, this specifies the address in the data buffer relative to the start of the buffer at which to start writing. When Mode = 1, this specifies the absolute address in processor memory at which to start writing.	

---

**Allocation Length** When Mode = 0, 4 bytes of header are transferred and then (*Allocation Length* - 4) bytes of data are transferred to the data buffer or mini buffer starting at address 0. When Mode = 1, Allocation Length bytes of data are transferred to processor memory. When Mode = 2, Allocation Length bytes of data are transferred to the data buffer or mini buffer. Note that when writing to the data buffer or mini buffer, it is possible for the combined length and offset to exceed the size of the buffer. The HP C1533A will wrap around from the buffer end and place the remainder of the Allocation Length at the start of the buffer.  
If Mode = 4, Allocation Length specifies the size of the download.

**Valid Field Values:**

The following table shows the values of the Buffer ID, Buffer Offset and Allocation Length fields for each valid value of Mode:

Mode	Buffer ID	Buffer Offset	Allocation Length
If Mode = 0 then:	0	0	≤ 8 00 00h+4
	2	0	≤ 2 00 00h+4
If Mode = 1 then:	1	Within valid controller processor address range	So that the window of processor memory starting at Buffer Offset is valid
If Mode = 2 then:	0	< 8 00 00h	≤ 10 00 00h
	2	< 2 00 00h	≤ 2 00 00h
If Mode = 4 then:	1	0	size of download

# WRITE FILEMARKS

10h

## Description:

WRITE FILEMARKS causes the specified number of filemarks or setmarks to be written beginning at the current logical position on tape.

If zero filemarks are to be written, the Immed bit must be zero. The drive writes any buffered data and marks to tape before reporting. This is the recommended way for a host to flush the buffer.

If setmarks are to be written, and the sum of the number to be written and the number already written exceeds 65,536, CHECK CONDITION is reported.

## Pre-Execution Checks:

Illegal Field Unit Attention	Flag Link Media Access	Bad LUN Media Write	Reservation Diagnostic Status	Deferred Error
---------------------------------	---------------------------	------------------------	----------------------------------	----------------

## Command Descriptor Block:

		Bit								
		7	6	5	4	3	2	1	0	
Byte	0	Operation Code (10h)								
	1	Logical Unit Number (0)				Reserved (0)			WSmk	Immed
	2	(MSB)								
	3	Number of Marks								
	4									
	5	Reserved (0)							Flag	Link
		(LSB)								

## CDB Fields:

WSmk	0	Filemarks are written
	1	Setmarks are written
Immed	0	Status will not be returned until the operation is complete. Note that the MODE SELECT Buffered Mode field must be zero for this.
	1	The drive returns GOOD status as soon as the Command Descriptor Block has been validated.
Number of Marks	The number of consecutive marks to be written to tape. Zero is only valid when immediate report is disabled.	

## Write Filemarks Specific Status:

Event	Status	Key	Notes
EOP/M encountered	CHECK COND'N	MEDIUM ERROR	Additional sense will be set to 0002h (EOP/M detected). The logical position will be at EOD. Subsequent unsolicited REQUEST SENSE commands will give tape position as EOD. It may still be possible to write a smaller quantity of information at this point (for example, filemarks, setmarks or a smaller block). This is because the logical position after a failure to write a block returns to the start of the unwritten block.
Early warning end-of-medium (EW) encountered	CHECK COND'N	NO SENSE	This is only reported if the write has been successful. Additional sense is set to 0002h (EOM/P detected). Residue information is valid and indicates zero blocks or bytes.
Completes properly and ends adjacent to a mark	GOOD	NO SENSE	The Mark bit is set. Additional sense will be set to 0001h (filemark detected) or 0003h (setmark detected).

### Residue Information:

Residue information depends on two variables:

- Fixed or Variable block mode.
- Immediate or Non-Immediate Report (Buffered mode).

Block Mode	Buffered Mode	Residue
Variable	0	Number of unwritten bytes
Fixed	0	Number of unwritten blocks
Variable	not 0	Total number of unwritten bytes and marks (including those buffered before the command was received). The size can be greater than the command operation size.
	specifically 1	Contains residues of all initiators with data in the buffer.
Fixed	not 0	Total number of unwritten blocks and marks (including those buffered before the command was received). The size can be greater than the command operation size.
	specifically 1	Contains residues of all initiators with data in the buffer.

Residue information is set in the sense data byte, with the Valid bit set and the information bytes set to the residue.





---

**Theory of Operations**

---

# Theory of Operations Contents

<b>TO1</b>	<b>Theory of Operations</b>	
	<b>Mechanism</b>	<b>1</b>
	Helical Scan Recording	1
	Four-Head design for Read-After-Write	2
	Motors and Sensors	2
	Motors	2
	Sensors	3
	<b>Electronics</b>	<b>4</b>
	Main Circuit Board	5
	Controller Electronics	5
	Device Electronics	5
	RF Amplification	5
	<b>Typical Signal Path</b>	<b>6</b>

The HP C1533A uses a tape mechanism which is developed from the principles used in digital audio tape (DAT) recorders, and uses the helical scan recording technique. Data is recorded according to the ECMA DDS standard, and communication with the computer system is by SCSI. This chapter introduces the mechanical functions of the tape drive and provides an overview of the electronics which link SCSI with the DAT mechanism.

## Mechanism


The mechanism used in the HP C1533A is similar to that used in a high quality digital audio tape recorder, however the mechanism used has been specifically designed to suit the special requirements of data storage.

### Helical Scan Recording

Most current computer tape drives and typical audio mechanisms (such as compact cassette) record on tracks along the length of the tape. Because of various limitations such as mechanical tolerance and magnetic crosstalk (interaction between signals from different tracks), it is very difficult to increase the data density on these products while retaining data integrity and compatibility.

Audio DAT overcomes this limitation, by recording tracks diagonally across the tape. This is achieved by mounting two heads on a rotating drum with an axis at  $6^\circ$  from the vertical. The drum rotates at 2000 rpm while the tape moves slowly (8 mm/sec) in the same direction. As a result, the heads, which are diametrically opposite, describe portions of a helix on the tape—hence helical scan. See figure TO-1.


Figure TO-1



Helical Scan  
Recording

In the HP C1533A, the drum rotates at 5737 rpm and the tape moves 15.5 mm/sec to achieve a higher transfer rate while writing narrower tracks for the DDS-2 format. Each head writes a track of data on the tape from bottom to top. The heads are wider than the tracks, so the tracks overlap with no wasted space between them. Crosstalk between the tracks is minimized by each head writing its data in angled strips along the track. The angle is called the azimuth angle. Each head is set with a different azimuth angle, so alternate tracks on the tape have their data written at different angles, as shown in figure TO-2.

Figure TO-2




Azimuth  
Angle

### Four-Head design for Read-After-Write

In the DDS-format drive design, two more heads are added to the rotating drum, making four at  $90^\circ$  to each other as shown in figure TO-3. These enable the drive to read data immediately after it has been written. If an error occurs, the drive can rewrite the erroneous frame repeatedly until it is read back successfully.

Figure TO-3



Four-Head  
Design

### Motors and Sensors

The locations of the motors, solenoids and sensors are illustrated in figure TO-4. The numbers in this section in italic type refer to items in figure TO-4.

#### Motors

The HP C1533A has four motors, which drive the following:

- The Head Drum *1*
- The Capstan Roller *4*
- The Supply Reel (S-Reel) *10*
- The Take-up Reel (T-Reel) *13*


They are direct-drive brushless motors which are quiet, reliable and very compact.

There is also a dc motor *2*, which controls the following operations:

- Engagement and disengagement of the Capstan Pinch Roller
- Tape threading
- Loading and unloading of the cartridge

This motor is a high-torque, brush-type motor.

Figure TO-4



## Sensors

The HP C1533A has eight mechanical switches which provide the following information:

- Whether a cartridge is present 14
- The open or closed state of six recognition holes on the cartridge 6 and 7

Combinations of open and closed holes give the tape drive with the following information:

- What type of tape is in the cartridge (e.g. Metal Particle)
- What thickness the tape is
- What the length of the tape is
- Whether the cartridge is write protected
- Whether it is a cleaning cartridge

An optical sensor 5 is used to identify BOM (beginning-of-media) because DAT tapes have partially transparent leaders. Audio and unqualified DAT tapes have fully transparent leaders, while DDS Media Recognition System qualified tapes have stripes on a transparent base. This allows the drive to recognize DDS-quality tapes.

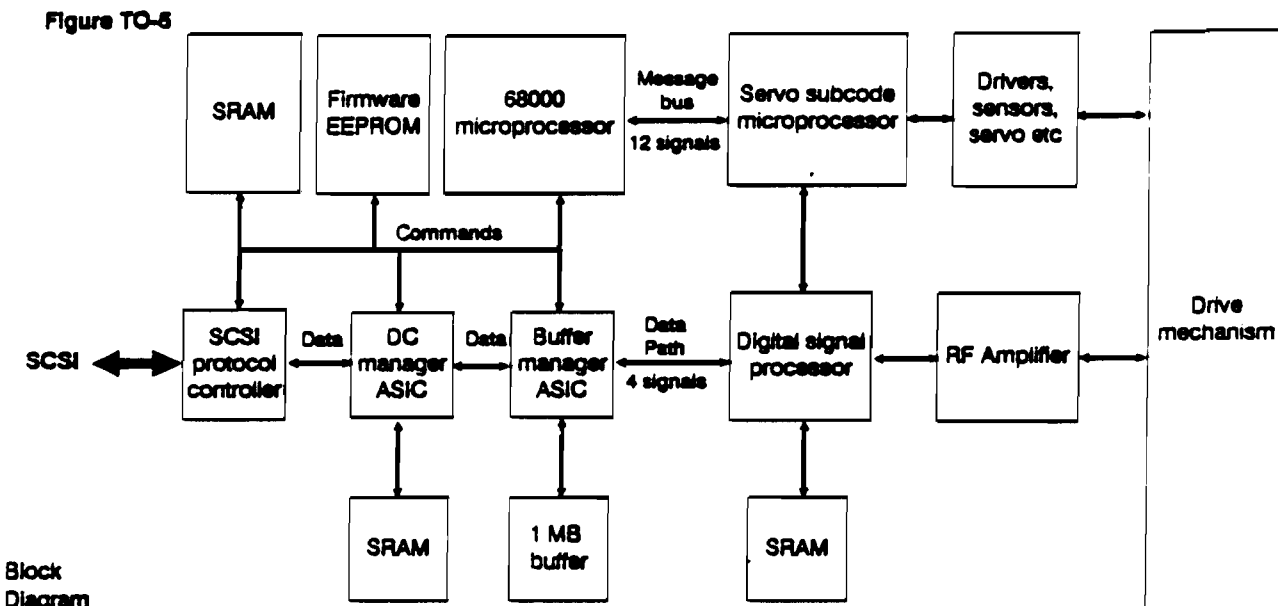
Magneto-resistive element 8 monitors the capstan speed.

Hall effect sensors built into the motors are used to monitor drum position 9, take-up reel speed 12 and supply reel speed 11.

Engagement or disengagement of the capstan pinch roller, tape threading and tape loading are all controlled by the mechanism mode cams. The mode cams are driven by the dc motor 2 mentioned earlier. Attached to the top of the mode cams is a three wiper wafer switch 3 which monitors the position of the cams, and consequently the mode of the drive.

## Electronics

A block diagram of the electronics which link the SCSI interface to the drive mechanism is illustrated in Figure TO-5.



The electronics consist of a main circuit board with components mounted on both sides. To reduce electro-magnetic interference the RF amplification circuitry is mounted on an additional small board. Signals are received through the SCSI interface and are interpreted as tape movement commands or data by the SCSI protocol controller. Data is treated by various encoding techniques to reduce the possibility of recovery errors and to make it compatible with the DDS format. These techniques include:

- Group Indexing, which allows the format to map variable length records into fixed group sizes.
- Checksum Generation, where the sum of a series of bytes is written to the tape, so that the figure can be checked against the sum of the same series of bytes when the tape is read.
- Randomization. The error rate for worst case data differs from random data by a factor of 10. Randomization reduces the worst case error rate by providing a stream of data which has a more consistent RF envelope.

Other duties performed by the electronics on the main circuit board are as follows:

- Tape management tasks. These include the following:
  - Mechanism Control
  - ATF (Automatic Track Following)
  - Servo Control
  - Diagnostics
  - Buffer management
  - Error detection
  - Control of the front panel display
  - Compression of data blocks

## Main Circuit Board

The main circuit board contains the following circuitry.

### Controller Electronics

#### *Controller Application Specific Integrated Circuit (ASIC)*

The Controller implements all DDS formatting and error correction functions on a single ASIC, as well as managing access to the 1 megabyte main data buffer.

#### *Buffer*

The buffer consists of four 1-megabit dynamic RAMs.

#### *Microprocessor*

The microprocessor is a 16 MHz 68000.

#### *Firmware*

The firmware is contained in two "Flash" EEPROMS which can be reprogrammed through the SCSI interface, or by a direct connection with the microprocessor bus at the factory. The firmware has six main components:

- A SCSI task which recognizes SCSI commands, messages and the SCSI phase protocols
- A DC task which controls the operation of data compression
- A buffer task which manages the buffering of data and the DDS format
- A device task which manages the reading and writing of data and controls the action of the drive mechanism through the Servo Microcomputer Interface
- A panel task which handles the front panel display and its operation
- An operating system which performs the self-test and schedules the tasks

#### *SCSI Protocol Controller*

Tape drives designed for single-ended operation use an HP ASIC (Application Specific Integrated Circuit) which has built-in, single-ended line drivers.

#### *Data Compression Manager ASIC*

Data compression and decompression is performed by an HP Compression/Decompression ASIC.

### Device Electronics

The Device Electronics perform the following functions:

- Converting of Servo Microcomputer Interface commands sent from the controller electronics into commands which cause movements of the tape drive mechanism
- Taking data from the buffer and writing it to tape
- Recovering data from the tape
- Separating and synchronizing the clock signal

The major component of the device electronics is a single chip microcontroller. It has a Hitachi H8 processor core, and contains many support functions including eight channel analog to digital conversion and facilities for automatic track following (ATF).

### RF Amplification

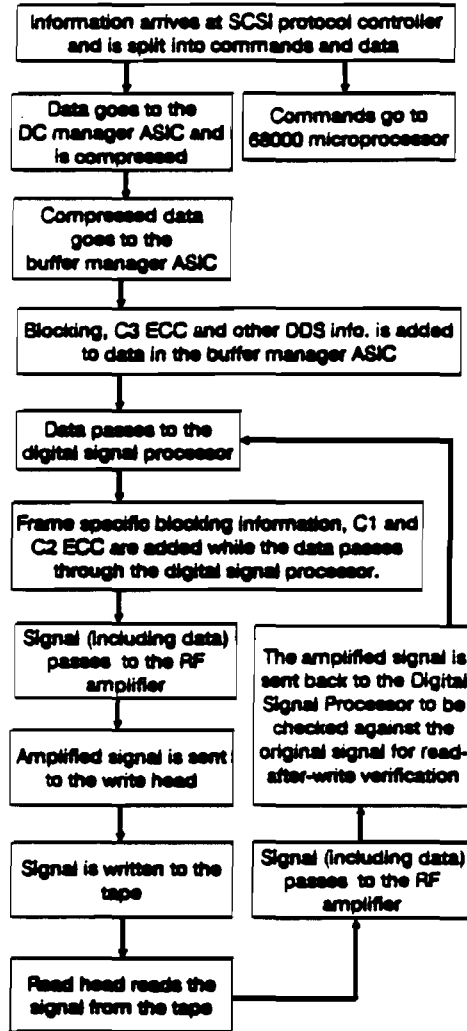
Data from the tape is amplified by an RF amplification stage which is contained on a separate board.

# Typical Signal Path

Figure TO-6 shows the path of a signal from the SCSI to the tape when the drive is writing data.

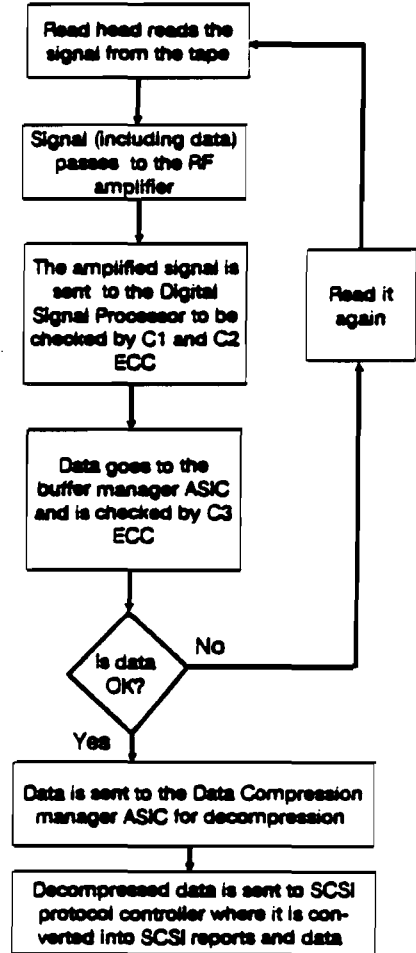
Figure TO-7 shows the path of a signal when it is read from the tape and passes through to SCSI.

Figure TO-6



Typical signal path:  
SCSI to tape

Figure TO-7



Typical signal path:  
Tape to SCSI



---

---

**Diagnostics**

---

# Diagnostics Contents

D1	Diagnostics Overview	
	Introduction	1
	Running a Test	2
	Invoking a Diagnostic Test from the Host	2
	Obtaining Test Results	3
	The Results Message	3
	Diagnostic Sequences	4
D2	Diagnostic Tests	
	The Power-Up Self-Test	5
	Host Action if the Power-Up Self-Test Fails	7
	Action to Take if a Test Fails	7

This section gives details of the diagnostic capabilities of the HP C1533A.

This chapter provides an overview and details of how to run tests and obtain results.

Chapter D2 covers the Power-Up Self-Test and the Self-Test Diagnostic Sequence which perform a variety of individual tests of the hardware and firmware.

The main purpose of the diagnostic firmware is to provide the user with a high level of confidence in the correct operation of the drive. If all of the diagnostics work successfully, the user should be confident that there is no problem with the drive.

## Introduction

For diagnostic purposes, the HP C1533A may be thought of as three distinct parts:

- The drive mechanism and associated electronics
- The media
- The external power supply

If a diagnostic test fails, a number indicates the most likely part where the failure occurred.

<u>Part</u>	<u>Number</u>
No problem found	0
Drive Mechanism and Associated Electronics	1
Media	2
External Power Supply	3

## Running a Test

Diagnostic tests can be run through the Power-Up Self-Test (see the next chapter) or through the host.

### Invoking a Diagnostic Test from the Host

The SCSI initiator (host) can instruct the drive to run a diagnostic test by sending the necessary parameters through the SCSI SEND DIAGNOSTIC command. The command has a six-byte Command Descriptor Block as follows:

		Bit								
		7	6	5	4	3	2	1	0	
Byte	0	Operation code (1Dh)								
	1	Logical Unit Number (0)			PF (0)	Reserved	Self-Test	Dev Off(0)	Unit Off.	
	2	Reserved								
	3	(MSB)	Parameter List Length						(LSB)	
	4									
	5	Reserved						Flag	Link	

The Flag and Link bits operate as described in Chapter 11 "Control Byte—Flag and Link Bits". Following the command descriptor block, eight bytes of data should be sent during the DATA OUT phase, as follows:

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Test Number (00h)							
	1	SOE	CSE	PT	Loop Count Identifier				
	2	Parameter A							
	3	Parameter B							
	4	Parameter C							
	5	Parameter D							
	6	Parameter E							
	7	Parameter F							

<b>Test Number</b>	The number of the diagnostic test to be performed.
<b>SOE</b>	The SOE (Stop On Error) bit determines how the loop count should be terminated if an error occurs. It only has meaning if the test is to be executed more than once. <ul style="list-style-type: none"> <li>1 Stop test immediately an error is detected</li> <li>0 Log the error and continue testing</li> </ul>
<b>CSE</b>	The CSE (Continue Sequence on Error) bit is similar in effect to the SOE bit, in that it is used during a sequence of tests to specify what should happen when an error occurs. The host may specify a sequence of commands to be executed, and this bit determines whether a failure of one test should halt the sequence, or allow it to continue. It only has meaning if the test is to be executed more than once. <ul style="list-style-type: none"> <li>0 Log the error and abandon the test sequence</li> <li>1 Log the error and continue the test sequence</li> </ul>
<b>PT</b>	The PT (Pass Through) bit tells the drive what to do when the diagnostic command has finished.

- 0 Return control to the normal firmware, so that the mechanism is returned to where it was before the test. However, if there was a tape in the mechanism, and as a result of the tests it is no longer present, the tape will stay in the ejected state.
- 1 Stay in "pass-through" mode; the mechanism will stay where the test left it. Pass-through mode will continue until a diagnostic command with PT = 0 is sent. Note that the drive will not be able to perform normal operations if it is left in pass-through state.

**Loop Count Identifier** The Loop Count Identifier determines how many times the test will be repeated before reporting back to the host, according to the following code:

- 0 continuous (it terminates only on an error or a power-cycle)
- 1 run once
- 2 run 10 times
- 3 run 100 times
- 4 run 1000 times

If the test completes successfully, GOOD status is returned in the Status phase of the SCSI sequence.

If the test fails, a CHECK CONDITION status is sent with a HARDWARE ERROR sense key and a DIAGNOSTIC FAIL additional sense key set. The host should then send a RECEIVE DIAGNOSTIC RESULTS command which will return information as to which test failed and the nature of the failure. Note that unless the SOE bit is 0, a failure will stop the test without completing the loop count.

**Parameters A-F** These are test specific; six parameters are not needed for most of the tests. Any unused parameters should be set to 00.

## Obtaining Test Results

The SCSI RECEIVE DIAGNOSTIC RESULTS command tells the drive to return the results of a diagnostic test to the host. If CHECK CONDITION status is returned on completion of a SEND DIAGNOSTIC command, together with the HARDWARE ERROR sense key and an Additional Sense Code of DIAGNOSTIC FAILURE, the host should issue a RECEIVE DIAGNOSTIC RESULTS command to obtain the data (indicating the actual failure and the Most Suspect Sub-Assembly). The messages are in the format of the Results Messages described in the next section.

The command has a six-byte Command Descriptor Block as follows:

		Bit									
		7	6	5	4	3	2	1	0		
Byte	0	Operation Code (1Ch)									
	1	Logical Unit Number (0)				Reserved					
	2	Reserved									
	3	(MSB)		Allocation Length							
	4										
	5	Reserved							Flag	(LSB) Link	

The Flag and Link bits operate as described in Chapter 11 "Control Byte—Flag and Link Bits".

## The Results Message

The results are returned as sixty-four bytes of data.

		Bit								
		7	6	5	4	3	2	1	0	
Byte	0	Test Number								
	1	Error Code								
	2	Most Suspect Sub-Assembly								
	3	Loop Count Number								
	4	(MSB)	Test Specific Information							(LSB)
	63									

<b>Error Code</b>	0	No error. The test has completed successfully.
	<i>n</i>	The test has experienced an error described by this error code.
<b>Most Suspect Sub-Assembly</b>	The part of the drive most likely to be at fault:	
	0	No error occurred, or no FRU could be distinguished
	1	Drive (Controller PCA and mechanism)
	2	Tape
	3	Power Supply
<b>Loop Count Number</b>	The number of loops that were completed before the test finished. Any number greater than 255 is reported as 255.	
<b>Test Specific Information</b>	See descriptions of the individual tests for details of this.	

If the command completes successfully, a status of GOOD is returned in the Status phase of the SCSI sequence. If the command fails, CHECK CONDITION status is returned. A REQUEST SENSE command should then be sent. The Sense Key and Additional Sense Codes returned by this command reflect the nature of the failure.

## Diagnostic Sequences

The HP C1533A will execute diagnostic sequences sent from the host. To do this, the host repeats the SEND DIAGNOSTIC parameter list, modifying the list for each diagnostic test required.

Once the drive has received the complete list, it can be left to execute each diagnostic in turn without further host intervention and report when the sequence is complete, or when an error occurs.

Because the tests run without host intervention, no reports are available apart from the final test executed. This will either be the last test on the list, or when the sequence is terminated by an error. The results are obtained using the RECEIVE DIAGNOSTICS RESULTS command, as described earlier.

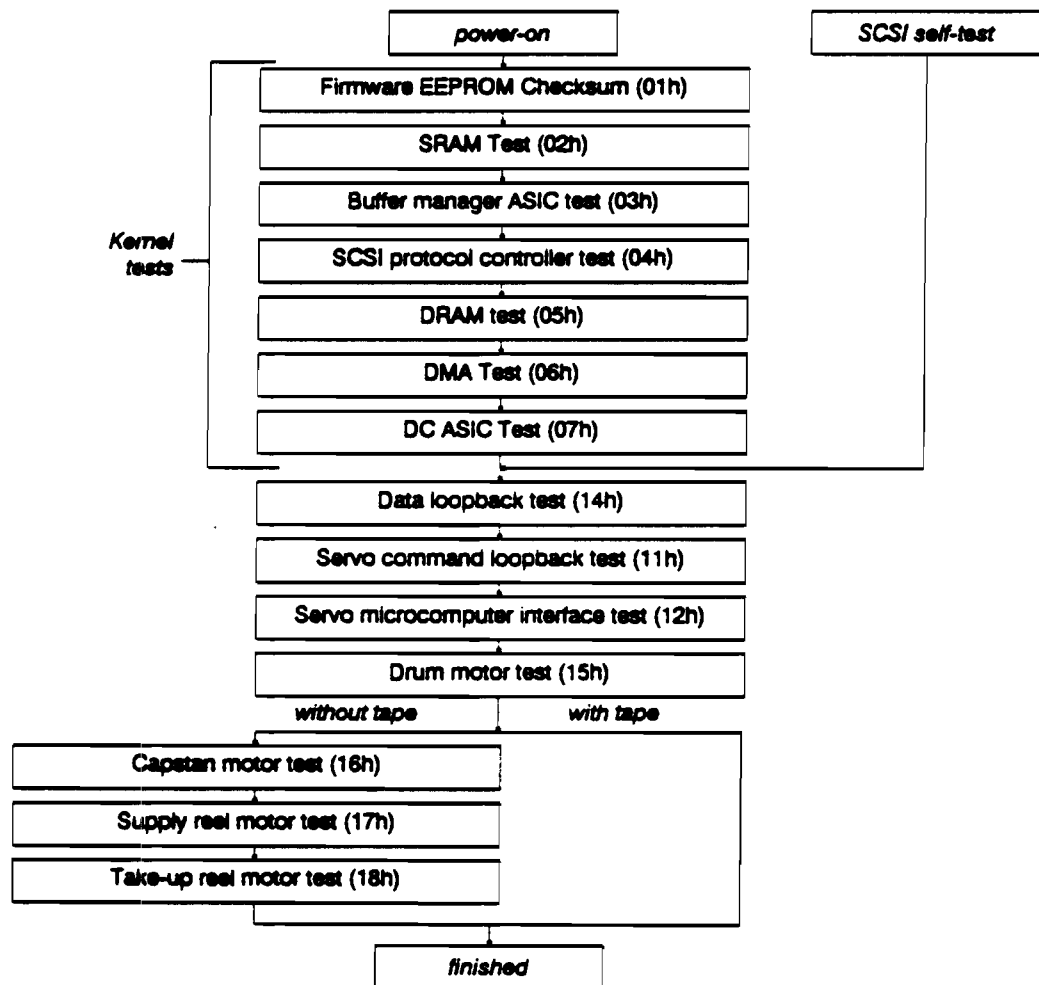
# Diagnostic Tests

## D2

This section covers the Power-Up Self-Test and the Self-Test Diagnostic Sequence (00h), which run a series of tests of the hardware and firmware.

When the drive is power-cycled, a series of tests are run to check that the drive has enough functionality to work correctly. The actual tests performed depend on whether a tape is present in the drive (this does not include a tape in the ejected position). The sequence is shown in the following diagram.

If the host requests a self-test through the SEND DIAGNOSTIC command with the Self-Test bit set to 1, the same sequence of tests is run, starting after the Kernel diagnostics (tests 01h through 10h).



Results of the tests are retrieved by the RECEIVE DIAGNOSTIC RESULTS command, described in Chapter D1. The tests are described below, together with the hexadecimal number of each test. This number is referred to in the Results Message if the sequence of tests failed.

"The drive hangs" means that it does not come online to SCSI.

Test Name	Description	Test Specific Results																
01h EEPROM checksum	Checks content of EEPROM by performing a word-wide checksum	If the test fails, the drive hangs.																
02h Processor SRAM Test	Non-destructive test of the processor SRAM. Includes writing walking 1s and walking 0s patterns over the entire SRAM area	If the test fails, the drive hangs.																
03h Buffer Manager ASIC test	Checks the operation and registers of the Buffer Manager Application Specific Integrated Circuit used to perform data transformations and hardware control tasks. It does not check for cross-register corruption, nor registers relating to the Servo microcomputer interface	If the test fails, the drive hangs.																
04h SCSI test	This checks the functionality of the SCSI interface ASIC by performing register checks, and by using its internal loopback mode.	If the test fails, the drive hangs.																
05h DRAM test	Verifies basic operation of DRAM buffer. Walking 1s and 0s patterns are used in a few locations and the RAM area around the base of each 1Kbyte page is checked	If the test fails, the drive hangs. Note: the host can check the buffer by using WRITE BUFFER and READ BUFFER commands.																
06h DMA test	Performs internal data transfers between the buffer and the SCSI interface, checking the data transferred	If the test fails, the drive hangs.																
07h DC ASIC test	Tests the Application Specific Integrated Circuit which controls data compression.																	
11h Servo micro-computer interface loopback test	Tests the communication path between the controller and the mechanism processor. The subcode area of the message buffer (55 bytes) is filled with data by the controller. The processor then reads the bytes, inverts each one and replaces the inverted bytes in the buffer. Finally the controller reads them back and checks that they are the inverse of the original data.	<p>If the comparison fails, the test specific bytes contain a copy of the bytes returned in the subcode area, and the final (64th) byte indicates the offset within the 55 bytes of the first byte found to be in error.</p> <table border="1"> <thead> <tr> <th>Byte</th> <th>Contents</th> </tr> </thead> <tbody> <tr> <td>5</td> <td>Number of bytes sent</td> </tr> <tr> <td>6</td> <td>1st data byte reported</td> </tr> <tr> <td>7</td> <td>2nd data byte reported</td> </tr> <tr> <td>:</td> <td>:</td> </tr> <tr> <td>n+6</td> <td>nth data byte reported</td> </tr> <tr> <td>:</td> <td>:</td> </tr> <tr> <td>64</td> <td>Offset of first byte in error</td> </tr> </tbody> </table>	Byte	Contents	5	Number of bytes sent	6	1st data byte reported	7	2nd data byte reported	:	:	n+6	nth data byte reported	:	:	64	Offset of first byte in error
Byte	Contents																	
5	Number of bytes sent																	
6	1st data byte reported																	
7	2nd data byte reported																	
:	:																	
n+6	nth data byte reported																	
:	:																	
64	Offset of first byte in error																	
12h Servo microcontroller self-test	The controller instructs the servo microcontroller to perform an internal self-test. The test includes performing a checksum of the ROM and hardware interface checks. It also checks that the firmware revision level is acceptable.	None.																
14h Data loopback test	Full, three-frame data loopback test. Data is transferred from the buffer manager ASIC across the servo microcomputer interface data bus into the signal processor, then back again.	<table border="1"> <thead> <tr> <th>Byte</th> <th>Contents</th> </tr> </thead> <tbody> <tr> <td>5 (MSB)</td> <td>Number of blocks sent (LSB)</td> </tr> <tr> <td>8</td> <td></td> </tr> <tr> <td>9 (MSB)</td> <td>Number of blocks in error (LSB)</td> </tr> <tr> <td>12</td> <td></td> </tr> </tbody> </table>	Byte	Contents	5 (MSB)	Number of blocks sent (LSB)	8		9 (MSB)	Number of blocks in error (LSB)	12							
Byte	Contents																	
5 (MSB)	Number of blocks sent (LSB)																	
8																		
9 (MSB)	Number of blocks in error (LSB)																	
12																		



15h	Drum motor test	Tests operation of drum, sensors and drivers. Drum is accelerated to normal speed then stopped. The time taken to accelerate to normal speed is checked.	Byte	Contents
			5 (MSB)	Actual time to get up to speed (LSB)
			6	to speed (LSB)
			7 (MSB)	Maximum acceptable time (LSB)
			8	acceptable time (LSB)
			Time is recorded in milliseconds. If the drum never got up to speed, the actual time taken is recorded as FFFFh.	
16h	Capstan motor test	Tests operation of capstan, sensors and drivers. The capstan is accelerated to a given speed then stopped. The time taken to accelerate is checked.	Byte	Contents
			5 (MSB)	Actual time to reach forward speed (LSB)
			6	forward speed (LSB)
			7 (MSB)	Max acceptable time to reach forward spd (LSB)
			8	to reach forward spd (LSB)
			9 (MSB)	Actual time to reach reverse speed (LSB)
			10	reverse speed (LSB)
			11 (MSB)	Max acceptable time to reach reverse spd (LSB)
			12	to reach reverse spd (LSB)
			Time is recorded in milliseconds. If the capstan never got up to speed, the actual time taken is recorded as FFFFh.	
17h	Supply reel motor test	Tests operation of supply reel, sensors and drivers. The motor is turned on, and after a fixed delay, the controller checks that it is rotating. The motor is then stopped. <i>Note: This test cannot be performed with a tape present.</i>	If the supply reel takes too long to reach the target speed, the test specific results area will contain the time taken for the acceleration.	
18h	Take-up reel motor test	Tests operation of take-up reel, sensors and drivers. The motor is turned on, and after a fixed delay, the controller checks that it is rotating. The motor is then stopped. <i>Note: This test cannot be performed with a tape present.</i>	If the take-up reel takes too long to reach the target speed, the test specific results area will contain the time taken for the acceleration.	

### Host Action if the Power-Up Self-Test Fails

If the Power-Up Self-Test fails, the drive returns a CHECK CONDITION status with a HARDWARE ERROR sense key set. In this state, the host is only permitted to unload the tape or execute diagnostics. The host should execute a REQUEST SENSE command to find out which self-test failed, the error code and the suspected sub-assembly. The display persists until the drive is reset.

### Action to Take if a Test Fails

The whole drive is a single Field Replaceable Unit (FRU). This means that if a diagnostic test fails, you should contact support.









**Special Features**

---

# Special Features Contents

<b>SF1</b>	<b>Supporting DDS Features</b>	
	<b>Partitioning</b>	<b>1</b>
	<b>Re-formatting</b>	<b>2</b>
	<b>Making or Modifying Two Partitions</b>	<b>2</b>
	<b>Determining the Number of Partitions</b>	<b>2</b>
	<b>Selecting the Active Partition (2-Partition Tapes)</b>	<b>3</b>
	<b>Determining the Current Partition</b>	<b>3</b>
	<b>Suggested Uses</b>	<b>3</b>
	<b>Check Conditions during Normal Operation</b>	<b>3</b>
	<b>Early Warning</b>	<b>4</b>
	<b>EOP/M</b>	<b>4</b>
	<b>Managing Blocks and Tape Marks</b>	<b>4</b>
	<b>Writing Fixed-Length or Variable-Length Blocks</b>	<b>4</b>
	<b>Writing Filemarks and Setmarks</b>	<b>5</b>
	<b>EOD Handling</b>	<b>5</b>
	<b>Finding Blocks, Filemarks, Setmarks and EOD</b>	<b>6</b>
	<b>Fast-Search</b>	<b>8</b>
	<b>Implementing Fast-Search</b>	<b>8</b>
	<b>Media Recognition System</b>	<b>8</b>
	<b>Media Recognition Active</b>	<b>8</b>
	<b>Media Recognition Disabled</b>	<b>8</b>
	<b>Error Handling</b>	<b>9</b>
	<b>Read-After-Write (RAW)</b>	<b>9</b>
	<b>C3 ECC</b>	<b>10</b>
	<b>N-Group Writing</b>	<b>10</b>
	<b>Verify</b>	<b>10</b>
<b>SF2</b>	<b>Software Decompression</b>	
	<b>Entities and Software Decompression</b>	<b>11</b>
	<b>Host/Drive Interaction for Software Decompression</b>	<b>11</b>
	<b>Summary</b>	<b>13</b>

This section provides information and advice for software engineers who are involved in developing applications for the HP C1533A DDS Tape Drive.

The HP C1533A is based on Digital Audio Tape (DAT) technology, and uses a recording format which has been specifically developed for computer applications.

In addition to standard backup applications, it is particularly suitable for the following:

- Archive and Retrieval
- Unattended Backup
- Data Interchange
- Software Distribution

This chapter deals with the SCSI-2 commands which are used to support the following:

- Partitioning
- Managing Blocks and Tape Marks
- Fast-Search
- Error Handling

Full details of all the supported SCSI commands can be found in the Interface section of this manual. Information is also provided on Check Conditions which may occur during normal operation of the drive.

Chapter SF2 provides information about Data Compression entities, and software decompression under host control.

---

## Partitioning

Creating two partitions allows a single tape to contain two separate areas which can be written and read independently. The size of each partition can be defined to suit application needs, and redefined later if necessary.

If a blank tape is loaded in the drive, it is used as a 1-Partition tape unless the host formats it as two partitions through the Medium Partitions Parameter page of the MODE SELECT command. The Medium Partitions Parameter page gives the host the following options:

- To generate two partitions on a blank tape
- To convert a 1-Partition tape to a 2-Partition tape
- To convert a 2-Partition tape to a 1-Partition tape
- To change the size of Partition 1 on a 2-Partition tape

---

**Note** When a tape is formatted, any data already on the tape is lost, including the Tape logs.

---

When formatting two partitions, the drive writes the Reference and System area of Partition 1, the Vendor Group, EOD frames up to Partition 1's EOP, the Reference and System area of Partition 0, the Vendor Group and a normal EOD area. All this is written as a continuum.

## Re-formatting

When a 2-Partition tape is re-formatted as 1-Partition, a new Reference and System area and Vendor Group are created at the front of the tape. In addition, a normal EOD Area, preceded by Data area amble frames, is written immediately after the Vendor Group.

## Making or Modifying Two Partitions

If you want a blank tape formatted as two partitions, or you want to change the number or size of partitions on an existing tape, follow these steps:

- 1 Determine the size required for Partition 1, which will depend on the intended usage.
- 2 Issue the SCSI MODE SELECT command, setting the appropriate values in the Medium Partition Parameters page (page code 11h) as follows:

		Bit							
		7	6	5	4	3	2	1	0
Byte	0	Page Code (11h)							
	1	Additional Page Length (08h)							
	2	Maximum Additional Partitions (01h)							
	3	Additional Partitions Defined							
	4	FDP (0)	SDP (0)	IDP (1)	PSUM (10)		Reserved (0)		
	5	Medium Format Recognition (03h)							
	6	Reserved (0)							
	7	Reserved (0)							
	8	(MSB)	Size of Partition 1						
9								(LSB)	

The minimum space allocated for Partition 1 is 1 megabyte. This allows sufficient space for the Reference and System area, Vendor Group, one data group before EW, the EW-to-EOP distance and the EOP-to-Boundary distance.

The page code 11h settings for converting a 2-Partition tape to 1-Partition are similar to those shown above, except that the Additional Partitions Defined field and the Size of Partition 1 field should both be set to 0. The tape will be initialized to contain one partition—Partition 0—spanning the whole tape.

## Determining the Number of Partitions

To determine the number of partitions, and the size of Partition 1 on a 2-Partition tape, send the MODE SENSE command with Page Code set to 11h, the Medium Partitions Parameter page. The number of partitions will be returned in the Additional Partitions Defined field (0 or 1). If the tape is 2-Partition, the size of Partition 1 (in megabytes) is returned in the Size of Partition 1 field.



## Selecting the Active Partition (2-Partition tapes)

When using a 2-Partition tape, the host must select the desired partition before issuing any read or write commands. This partition then becomes the Active Partition.

This is done by using the MODE SELECT command, and setting the appropriate values in the Device Configuration Page (page code 10h). The entries used are the Change Active Partition (CAP) bit and the Active Partition field. To select Partition 1, set the CAP bit to 1 and change the value of the Active Partition field to 10h.

Any values other than 00h or 01h in the Active Partition field will be rejected by the drive. Also, if the tape has not been formatted into two partitions, any attempt to change partitions will be rejected.

When a tape has just been loaded, the Active Partition field will be set to 00h regardless of how the tape has been formatted. The drive will also default to Partition 0 if there is a power cycle, BUS DEVICE RESET message or hard reset.

Once the Active Partition has been selected in this way, any subsequent read, write or tape positioning commands only apply to this partition.

## Determining the Current Active Partition

To determine the currently active partition on a 2-Partition tape, use the MODE SENSE command and select the Device Configuration Page (page code 10h). The currently active partition number will be returned in the Active Partition field.

## Suggested Uses

A possible use of the 2-Partition structure is for Partition 1 to contain a directory of the data files held in Partition 0. The directory could then be accessed very rapidly by the drive to get information about the location of data files. Such a directory could be written to Partition 1 after the data files have been written to Partition 0.

## Check Conditions during Normal Operation

The following section describes the error conditions which may occur during normal operation of the drive, using either a 1-Partition or a 2-Partition tape.

Table SF-1 provides a list of SCSI commands which will return a CHECK CONDITION whenever EOD, EW and EOP/M are detected by the drive when not expected. This list also applies if EW and EOP are detected for Partition 1 on 2-Partition tapes when not expected.

Table SF-1

	EOD	EW	EOP/M
	LOCATE	LOAD/UNLOAD	LOCATE
	READ	WRITE	READ
	SPACE	WRITE FILEMARKS	REQUEST SENSE
	VERIFY		SPACE
			VERIFY
SCSI Commands involving EOD, EW and EOP/M			WRITE
			WRITE FILEMARKS

**Note** For full details of all SCSI commands supported by the HP C1533A, refer to the Interface sections.

## Early Warning

The early warning mark (EW) is generated internally by the drive when there is approximately 500 mm of tape remaining before the End-of-Partition/Medium (EOP/M). EW will be signalled as a CHECK CONDITION by the commands listed in table SF-1.

Normally, when an EW warning is received by the host, it will stop sending data to the drive and issue a REWIND command. This will force the buffer to flush and write EOD. The 10 megabytes remaining are adequate to take the contents of the buffer and allow for any rewrites due to poor media. Note that it is up to the host to decide what action should be taken when the EW warning is generated.

---

**Note** EW warnings are not generated on read passes.

---

The EW used in the first partition of a 2-Partition tape is generated 500 mm before the EOP of Partition 1.

## EOP/M

Two conditions within the drive cause a EOP/M CHECK CONDITION:

- 1 When using a 2-Partition tape, it indicates EOP, the point in Partition 1 beyond which the drive will not allow data to be written within that partition.
- 2 When using a 1-Partition tape, or writing in the second partition (Partition 0) of a 2-Partition tape, it signals the end of the usable media.

See table SF-1 for a list of EOP/M handling commands.

---

## Managing Blocks and Tape Marks

This section describes the SCSI-2 commands which are used to handle blocks and tape marks. Full details of all SCSI commands supported by the HP C1533A can be found in the Interface sections.

### Writing Fixed-Length or Variable-Length Blocks

The MODE SELECT command (15h) is used to select either fixed-length or variable-length block mode of operation.

If fixed-length blocks are required, enter the desired block length in bytes in the Block Length field of the MODE SELECT Parameter Block Descriptor. If variable-length blocks are required, enter 0 in the Block Length field.

---

**Note** Selecting variable-length mode in this way is not mandatory. The READ, WRITE and VERIFY commands may force variable-length mode by setting the Fixed Bit to 0.

---

When issuing the READ, WRITE or VERIFY commands in fixed-length mode, the Fixed Bit field in the Command Descriptor Block (CDB) must be set to 1. The number of blocks to be transferred is entered in the Transfer Length field of the CDB. If the Fixed Bit is set to 1, but the Block Length value of the MODE SELECT command is 0, the command will be rejected with a CHECK CONDITION status and an ILLEGAL REQUEST sense key set.

When issuing READ, WRITE or VERIFY commands in variable-length mode, the Fixed Bit field in the CDB must be set to 0. The number of bytes to be transferred must be entered in the Transfer Length field. Note that it is permissible to select variable-length mode even if the Block Length value of the MODE SELECT command is not zero.

## Writing Filemarks and Setmarks

Both types of tape mark are created using the WRITE FILEMARKS command (opcode 10h).

If the WSmk bit in the CDB is 1, a setmark is written to tape.

If the WSmk bit in the CDB is 0, a filemark is written to tape.

Multiple filemarks or setmarks can be written by entering the number required in the Number of Marks field.

---

**Note** Details of the maximum number of setmarks, filemarks and blocks that can be written to a single tape are given in the section "How Indexing Works" of chapter DF2.

---

The setmark is considered hierarchically superior to the filemark. It allows the host software the freedom to include any number of blocks and filemarks within a set and be able to search to the end of it in one motion. Searching to setmarks automatically invokes the drive's fast-search capability.

The host software does not need to know the number of blocks or filemarks contained within a set in order to position past it.

When writing setmarks or filemarks, the timing for the command will vary considerably depending on the mode chosen. The performance within a chosen mode will vary depending on whether or not the Immediate bit is set.

### Buffered Mode—Immediate bit set

In this mode the filemarks are included in the data stream from the buffer to the tape, and streaming performance can be maintained.

### Buffered Mode—Immediate bit not set

If the Immediate bit is *not* set, the following sequence will occur:

- 1 Any separator marks are appended to the data in the buffer.
- 2 The buffer is then flushed to tape. GOOD status will not be returned until the tape mark has been successfully written using the read-after-write capability.

## EOD Handling

Marking the end of the recorded data serves two purposes:

- 1 It limits fast-searching to the area of the tape containing recorded data, and no time is wasted searching unrecorded areas.
- 2 It shows which data is valid. When data is written to the tape it may either be appended to the tape's existing contents, starting at the current EOD, or it may be written over existing data. If no overwriting is done, the Data area grows and the EOD area is written nearer the end of the tape. If existing data is overwritten, and there is less new data than was originally on the tape, the EOD area is written closer to the beginning of the tape. In this case, any recorded data beyond the new EOD area is no longer valid.

The drive flushes the write buffer to tape and appends an EOD marker after receipt of the following non-write commands:

LOAD/UNLOAD	MODE SELECT	READ BUFFER	SEND DIAGNOSTIC	VERIFY
LOG SENSE	READ	REWIND	SPACE	WRITE BUFFER

The buffer can be maintained through the following media commands, assuming that no other flush condition has been met—for example, write hold-off timeout:

ERASE                      WRITE FILEMARKS (Immediate)                      WRITE

The tape can be positioned at EOD by using the SPACE command (see "Finding Blocks, Filemarks, Setmarks and EOD" later in this chapter). It follows that the convention of using two filemarks to signify end of data need not be used with DDS drives, though spacing to sequential filemarks is supported. If EOD is detected by the drive when not expected, a CHECK CONDITION will be returned for the commands in table SF-1.

## Finding Blocks, Filemarks, Setmarks and EOD

The SCSI SPACE command (opcode 11h) is used for positioning the drive in relation to the blocks and tape marks on the tape. A 3-byte code is entered in the Code field of the CDB to indicate what is to be spaced over. The code is as follows:

DB1(2)	DB1(1)	DB1(0)	Description
0	0	0	Blocks
0	0	1	Filemarks
0	1	0	Sequential Filemarks
0	1	1	End-of-Data (EOD)
1	0	0	Setmarks
1	0	1	Sequential Setmarks

When spacing over blocks or marks, the Count field of the CDB is used to determine the direction of movement and the number of items to be spaced over. It is interpreted as follows:

- A positive value of  $N$  causes forward movement (towards EOM) over  $N$  blocks or marks. The tape is logically positioned after the  $N$ th block or mark on the EOM side.
- A zero value causes no change in the logical position, except when spacing to EOD.
- A negative value  $-N$  (two's complement notation) causes reverse movement (towards BOM) over  $N$  blocks or marks. The tape is logically positioned on the BOM side of the  $N$ th block or mark.

## Implementation of the SPACE Command

Spacing to	Event Detected	Valid	Mark	EOM	Sense Fields		Note
					Sense Key	Additional Sense	
<b>Blocks</b>	Filemark	1	1	0	NO SENSE	0001h Filemark detected	a, b
	Setmark	1	1	0	NO SENSE	0003h Setmark detected	a, b, d
	EOD	1	0	0	BLANK CHECK	0005h EOD detected	b, e, h
	BOP	1	0	1	NO SENSE	0004h BOP/M detected	b, f
	EOP/M	1	0	1	MEDIUM ERROR	0004h BOP/M detected	b, g
<b>Filemarks</b>	Setmark	1	1	0	NO SENSE	0003h Setmark detected	a, b, d
	EOD	1	0	0	BLANK CHECK	0005h EOD detected	b, e, h
	BOP	1	0	1	NO SENSE	0004h BOP/M detected	b, f
	EOP/M	1	0	1	MEDIUM ERROR	0002h EOP/M detected	b, g
<b>Sequential Filemarks</b>	Setmark	0	1	0	NO SENSE	0003h Setmark detected	a, d
	EOD	0	0	0	BLANK CHECK	0005h EOD detected	e, h
	BOP	0	0	1	NO SENSE	0004h BOP/M detected	f
	EOP/M	0	0	1	MEDIUM ERROR	0002h EOP/M detected	g
<b>Setmarks</b>	EOD	1	0	0	BLANK CHECK	0005h EOD detected	b, e, h
	BOP	1	0	1	NO SENSE	0004h BOP/M detected	b, f
	EOP/M	1	0	1	MEDIUM ERROR	0002h EOP/M detected	b, g
<b>Sequential Setmarks</b>	EOD	0	0	0	BLANK CHECK		e, h
	BOP	0	0	1	NO SENSE	0004h BOP/M detected	f
	EOP/M	0	0	1	MEDIUM ERROR	0002h EOP/M detected	g
<b>EOD</b>	BOP	0	0	1	NO SENSE	0004h BOP/M detected	c, f
	EOP/M	0	0	1	MEDIUM ERROR	0002h EOP/M detected	c, g

**Note a** The final position will be located on the EOM side of the mark if the movement was forward (towards EOM), and on the BOM side of the mark if the movement was backwards (towards BOM).

**Note b** The Information field is set to the absolute value of the difference (residue) between the requested count and the actual number of blocks or marks spaced over.

**Note c** The Information field will contain no residue count and therefore the Valid bit is not set.

**Note d** The drive will only report that a setmark has been detected, while spacing over blocks or filemarks, if it has been configured through MODE SELECT to Report Setmarks (The RSmk bit set to 1 in the Active Format field of the Device Configuration page). Otherwise the drive will continue the space operation and the presence of the setmark will be transparent to the host.

**Note e** The tape is positioned such that a subsequent WRITE would append data after the last entity that has been written to the tape before EOD.

**Note f** The tape is physically positioned at BOP/M.

**Note g** The tape is physically positioned at EOM/P.

**Note h** If the drive meets EOD after the Early Warning (EW) mark, the EOM bit in the sense data is set.

Use of the SPACE command may or may not involve the drive's fast-search capability. The conditions for invoking fast-search are described after the next section.

---

## Fast-Search

The DDS format encodes information about the locations of separator marks in the Sub-Code areas of the tape. The Sub-Code area information can be read very rapidly, allowing the drive to find any block, filemark or setmark on the tape in an average time of 40 seconds for a 120 m tape. After the data group has been located, the data is read at normal speed. The decision to invoke fast-search is the drive's responsibility; it cannot be invoked by a command from the host system. See chapter DF2 for further details.

---

**Note** The timing of fast-search will depend on where the desired block is located on tape. When the drive is at BOP, the search time for blocks near the beginning of the tape which require fast-search will be approximately 7 seconds. For blocks near the end of the tape, the search time will be approximately 80 seconds.

---

## Implementing Fast-Search

Fast-search is only used under certain conditions. When a *SPACE* command is issued, and when the operation cannot be completed within the contents of the drive's 1 megabyte buffer, the drive will move the tape either by reading groups or by searching. The decision about which to use is made by an algorithm in the drive's firmware.

---

## Media Recognition System

The Media Recognition System enables the drive to identify when a DDS Media Recognition System cartridge is loaded in the drive, by detecting a striped area on the leader tape. If the leader tape has no striped area, the drive can be configured to react in the following ways:

- It can ignore the information.
- It can prevent writing to the cartridge, but allow it to be read.

These options are configured through switch 3 of the option switches on the underside of the drive:

Switch 3, State

off, Media recognition is active.

on, Media recognition is disabled. This is the default.

## Media Recognition Active

If the Media Recognition System is enabled, a DDS Media Recognition System cartridge will load and function in the normal way.

If a cartridge is inserted which is not DDS Media Recognition System, it is treated in the same way as if a write-protected cartridge had been inserted. An attempt to write to the tape results in a *CHECK CONDITION* with sense key of *DATA PROTECT* and an additional sense code of 3002h (incompatible media).

## Media Recognition Disabled

If the Media Recognition System is disabled, stripes on the leader tape are simply ignored, and tapes of any particular length are treated the same, whether DDS Media Recognition System tapes or not.

## Error Handling

The HP C1533A has three error-rate improvement features which can be selected by the host system. These are as follows:

- Read-After-Write (RAW)
- Third Level Error Correction (C3 ECC)
- N-Group Writing

By default, RAW and C3 ECC are enabled, and N-Group writing is disabled.

### Read-After-Write (RAW)


The DDS format supports a Read-After-Write technique to ensure that any data written to the tape can be read back without error. Each frame is examined after it has been written to check that it has been recorded correctly.

When a frame is identified as bad, it is rewritten later down the tape after a number of dummy frames have been written. Each frame, and the two in sequence after it, can be rewritten multiple times, thus having the effect of skipping over bad areas on the tape. The number of instances of a repeated sequence is set at a default value in the drive's firmware. When this number is exceeded, the writing of that group is aborted and the drive reports a hard error.

Also, in order to limit the physical group size, the maximum number of RAW frame rewrites that are allowed is set in the drive's firmware.

When reading a group, any rewritten frames need to be identified. If a frame has been rewritten, there will be more than one frame with the same Logical Frame Number. To recover the data in the group, all that is required is to make sure that at least one of each logical frame is read in the correct order. Any duplicates are ignored. If any frame is unreadable, the drive reads ahead up to eight frames to see if the frame has been rewritten. If it has, the rewritten frame is used and the drive continues reading from that point.

The following diagrams show examples of different frame rewrite conditions. The examples assume that two dummy frames are written. The numbers in the boxes represent the Logical Frame Numbers.



To perform Read-After-Write, two extra heads have been added to the drum, making a total of four at 90° to each other.

### Implementing RAW

RAW is switched on or off through the MODE SELECT command (opcode 15h) by changing the RAW bit (byte 2 bit 4) of the Device Configuration Parameters page (page code 10h).

If the RAW bit is 0, RAW is enabled—this is the default.

If the RAW bit is 1, RAW is disabled.

### C3 ECC

C3 ECC introduces an extra level of error correction in addition to those offered by the audio DAT format. This feature has the ability to correct any two tracks which are bad in a group. The error correction bits are stored in an additional (ECC) frame at the end of each group.

### Implementing C3 ECC

C3 ECC is switched on or off through the MODE SELECT command (opcode 15h) by changing the value of the C3 ECC bit (byte 2, bit 3) of the Device Configuration Parameters page (page code 10h).

If the C3 ECC bit is 1, C3 ECC is enabled—this is the default.

If the C3 ECC bit is 0, C3 ECC is disabled.

### N-Group Writing

N-Group Writing is a technique that writes every group more than once. Each group is repeated on the tape a number of times before the next group is written. Each set of rewritten groups is contiguous and contains instances of only one group.

The data and indexes of the rewritten groups are identical.

The maximum number of times a group can be repeated is 7; that is, up to 8 instances may be written.

### Implementing N-Group Writing

N-Group Writing is controlled through the MODE SELECT command (opcode 15h) by changing the N-Group field (byte 2, bits 0–2) of the Device Configuration Parameters page (page code 10h).

To enable N-Group Writing, enter a value from 001b to 111b in the N-Group field to define the number of times that each group must be repeated.

To disable N-Group Writing, enter 0 in the N-Group field. This is the default.

### Verify

The HP C1533A also supports the SCSI VERIFY command (opcode 13h), which provides an additional error-handling feature. This command verifies one or more blocks, beginning with the next block on the tape. This provides media verification only; no data is transferred between the host and the drive.



## Entities and Software Decompression

On a typical DDS-DC tape, two different data item types may be encountered:

- An uncompressed block
- An entity of compressed data

Entities are written by DDS-format drives which support on-board data compression. An entity is made up of a number of blocks of the same size compressed using the drive's compression algorithm. They are prefixed by an entity header which is an uncompressed descriptor containing information about the data within the entity.

A drive which supports compression algorithm *N* can decompress entities of type *N* and return the decompressed data to the host transparently. It can also return uncompressed records to the host if it finds them on the tape.

If this same drive finds an entity of type *M* (written by a different drive using a different compression algorithm), it will not usually be able to decompress it. Similarly, if a drive which does not support data compression at all finds an entity on tape, it cannot decompress it.

Some hosts may support software decompression; that is, the host computer is capable of decompressing entities. In these circumstances, the drive must be able to return a compressed entity to the host without decompressing it. The host must be aware that the data it is receiving is compressed and not a normal uncompressed record.

The following sections describe the use of the Data Compression Characteristics Page by DDS, DDS-DC, and DDS-2 drives. The page has code 0Fh in the MODE SELECT command.

### Interaction Between Host and Drive for Software Decompression

#### DC Drive—initial operating mode example

After drive reset, a drive which supports data compression can be configured to power-up with compression enabled and algorithm *N* selected. The Data Compression Characteristics Page contains the following values:

Field	Value	Meaning
DCE	1	Compression enabled
DCC	1	Compression capable device
DOE	1	Decompression enabled
RED	1	CHECK CONDITION returned on format change only
Compression Algorithm	<i>N</i>	Compression algorithm <i>N</i> enabled

### Non-DC Drive—initial operation mode example

A drive which does not support data compression will typically power-up with the Data Compression Characteristics Page containing the following values:

Field	Value	Meaning
DCE	0	Compression disabled
DCC	0	Non Compression capable device
DDE	0	Decompression disabled
RED	0	CHECK CONDITION returned on encountering compressed data
Compression Algorithm	0	No compression algorithm selected

### Example of Software Decompression Control

The host issues a number of READ commands to the drive, all of which are successful. On the next READ command, however, the drive detects an entity of type *M* on the tape, where *M* is an algorithm which the drive does not support. The drive cannot decompress the entity, but returns either the number of bytes in one block, or the total number of bytes in the entity, whichever is smaller. It is necessary for the drive to inform the host that it has found a data item on the tape which it cannot decompress. It does this by issuing a CHECK CONDITION to the host and setting the sense data as follows:

Field	Value	Meaning
Valid	1	Indicates that the Information field contains residual information from the failed READ command. Note that this will only be set if the entity length was different from the requested block length.
Sense Key	(00h)	NO SENSE. If encountered data is uncompressed.
	(01h)	RECOVERED ERROR. If encountered data is decompressable (and has been decompressed) by the drive.
	(03h)	MEDIUM ERROR. If encountered data is compressed and not decompressable by the drive.
Information	READ residue	The READ command failed with a residue as given in this field. Note that this will only be set if the entity length was different from the requested block length.
Command-Specific Information	Number of blocks in data item	The number of blocks in the entity is obtained from the entity header, which the drive can read. Note that in the case of a compressed-to-uncompressed format change, this field will contain 1 to indicate that one uncompressed record was encountered.
Additional Sense Code	(70h)	This ASC indicates the reason for the CHECK CONDITION as being a DECOMPRESSION EXCEPTION.
Additional Sense Code Qualifier	NN	Algorithm identifier for compressed entity encountered. Note that in the case of a compressed-to-uncompressed format change, this field will contain 0 to indicate that uncompressed data was encountered.

The drive is now positioned on the EOP side of the entity. If the host does not support software decompression, the drive can continue reading, if the host wishes. Note that in most cases, this type of host would clear the RED bit during initial device configuration, so that it would only ever receive a DECOMPRESSION EXCEPTION when the drive found an entity. Any blocks would be returned without this type of CHECK CONDITION.

If the host supports software decompression, it must check the sense data to see if it has received all the data from the entity. If it is reading in variable mode, it does this by looking at the residual count in the Information field. If this field is not negative, the host

has received all the compressed data and so will not need to SPACE reverse and re-read the entity.

On the other hand, if the Information field is negative, the requested block length is less than the actual entity length, and the host must SPACE reverse and re-read the whole entity in order to perform the software decompression successfully. The host does this by looking in the Command-Specific Information field in order to find the number of blocks in the entity. It then issues a SPACE reverse with the Count field set to the two's complement of this value. This will position the tape at the start of the entity. (Note that by subtracting the Information field—that is, the residual count—from the requested block length, the host can determine the actual entity size and reserve enough buffer space to receive the data.) Because the RED bit is set, the host will be able to READ the entity without the drive generating a CHECK CONDITION.

As long as the RED bit is set, the host will be able to continue reading entities from the drive until either an entity is found which has been compressed using a different algorithm, or uncompressed data is found.

Note that in fixed mode, the host will not be able to determine the size of the encountered entity from the Information field, because the residual information will be in terms of blocks, not bytes. It is up to the host in this case to take the appropriate action.

Note also that whenever the host needs to SPACE reverse over an entity because it has not managed to read all the data the first time round, the drive will return to its initial state as far as the RED bit is concerned. It will therefore issue a DECOMPRESSION EXCEPTION CHECK CONDITION in response to the following READ command.

## Summary

The example outlined above illustrates the requirements from both a host and a drive if software decompression is to be supported within a system. The interaction depends upon the drive being aware of compression/non-compression data boundaries and being able to report them, upon the host having the capability to perform software decompression, and upon the drive's ability to return compressed data items to the host without decompressing them.



---

---

**DDS Format**

---

## DDS Format Contents

<b>DF1</b>	<b>Digital Data Storage (DDS)</b>	
	DDS Format	1
	Tape Layout	2
	One-Partition Tape	2
	Device Area	3
	Reference and System Area	3
	Data Area	3
	EOD Area	4
	EW and EOP/M	4
	Tape Partitions	4
	Two-Partition Structure	5
	Entities (DDS-DC)	6
<b>DF2</b>	<b>Fast-Search, Indexing, Error Correction</b>	
	How Fast-Search Works	7
	How Indexing Works	8
	Block Access Table	9
	Example of Indexing	10
	How Error Correction Works in DDS	11
	Audio Format C1 and C2 ECC Performance	11
	DDS-Format C3 ECC Performance	12
	N-Group Writing Performance	14
	Read-After-Write	15
	Data Randomizer	16
	Checksum	16

## DDS Format


Digital Data Storage (DDS) is a recording format which supports the use of Digital Audio Tape (DAT) for computer applications. DDS tape drives make use of DAT features like helical-scan recording (see "Theory of Operations") and sophisticated error-correction techniques developed for use in the audio market. This chapter specifically details the second generation DDS-2 format, but most of it applies to DDS-1 and DDS-DC as well.

In addition to the error-correction and data-integrity features provided by DAT, the DDS format has the following features:

- A fast-search capability.
- The option of formatting the tape into a one-partition or two-partition structure.
- A third level of error correction (C3 ECC), which can recover errors that are too severe for the basic DAT format techniques (C1 and C2 ECC) to correct.
- A Read-After-Write (RAW) facility which checks the data for errors immediately after it is written, and rewrites it if necessary.
- The option of N-Group Writing (Multiple Group Writing), where each group of data is repeated a specified number of times before the next group is written.

The DDS format is structured to overlay the basic audio format. It does this by organizing the frames (pairs of tracks) of the audio format into a sequence of data groups on the tape, each group with a fixed data capacity. Figure DF-1 shows the construction of a group.


Figure DF-1



Construction of a Group

Each group consists of 22 frames-worth of data, where a frame is a pair of tracks across the tape. See figure DF-2.

Figure DF-2



One Frame

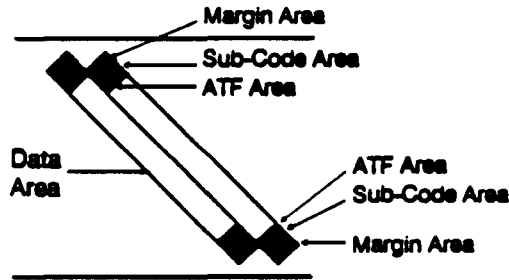
**Tape Layout**

As with the audio format, 60% of the length of a track is user data. The rest consists of:

- **Margin Areas**, which are used as guard bands.
- **Sub-Code areas**, which are used to enable fast-searching. This enables the drive to access any file on the tape rapidly, on average within 40 seconds for a 120 m tape. (See "Fast-Search" in chapter SF1 for more information).
- **ATF (Automatic Track Finding) areas**, which are used to center the head on the track.

Figure DF-3 shows the structure of a frame.

**Figure DF-3**



The Structure of a Frame

A host computer sends data and separator marks to a tape drive which supports the DDS format. The separator marks identify where logical collections of data (for example, files and sets of files) begin and end. The tape drive organizes the information into groups and writes it to tape. An index in each group identifies and locates the data blocks and separator marks contained in the group, and each group can be followed by an optional error-correction frame. (See "Error Handling" in chapter SF1 and "How Error Correction Works in DDS" in chapter DF2 for details of Error Correction techniques).

The method of indexing allows for fixed and variable length blocks and for separator marks to be encoded onto the tape without significantly affecting the amount of data that can be stored on a particular cartridge.


## Tape Layout

The DDS format supports the formatting of tapes into a 1-Partition or 2-Partition structure. These are described in the following sections.

### One-Partition Tape

The overall tape layout consists of four areas: the Device area, the Reference and System area, the Data area and the EOD area. Figure DF-4 shows the different areas.

**Figure DF-4**



Beginning-Of-Medium (BOM) and End-Of-Tape (EOT) are the points where the magnetic tape is joined to leader and trailer tapes respectively.

**Note** In the following, write operations can only occur if the tape is write-enabled.



## Device Area

This has three sections:

- 1 The *load section*, which is the part of the tape that is wrapped around the drum when the tape is first loaded.
- 2 A *test section*, where read and write tests of the drive's electronics and servo are performed.
- 3 A *guard section*, which provides a safety zone between the test section and the start of recorded frames.

## Reference and System Area

**Reference Area** The Reference area defines the Beginning-Of-Partition (BOP), and facilitates efficient positioning when updating the System area.

**System Area** The tape logs of usage and soft error occurrence are stored in the System area. This logged information is lost whenever a tape is formatted (see "Partitioning" in chapter SF1).

## Data Area

The Data area is written as a sequence of groups (see "DDS" earlier in this chapter), starting with a special Vendor Group which is written automatically by the drive. The Vendor Group holds details of the drive which created the partition, and the date. It is followed by data groups, which have a fixed capacity, and are used to store blocks, filemarks and setmarks written by the host. Fixed-length or variable-length data blocks may be written. These, together with tape marks, are mapped into the fixed capacity groups by a method known as Indexing, which uses a minimal amount of the data capacity of the tape (see chapter DF2).

In the DDS-DC format, used by the HP C1533A, data blocks are compressed into entities, which consist of a header and one or more compressed data blocks. See the section on Entities later in this chapter.

### Blocks

The host has the option of writing either fixed-length or variable-length blocks to the drive. The drive maps these blocks into the fixed-length group structure using the index in each group. A block may be written within a single group or span several groups, depending upon its size. The mapping process is invisible to the host system.

If the data block is compressed, along with others, into an entity, then the entity is mapped into the fixed-length group structure in a similar way. Again, this process is invisible to the host system.

### Separator Marks

DDS provides for two types of separator marks—filemarks and setmarks—which are each represented by 4 bytes in the index of a group. It is the responsibility of the application developer to define the logical significance of these marks.

**Filemarks** The meaning of filemarks is defined by the host.

**Setmarks** Setmarks provide an additional method of data segmentation. This new type of mark gives the host the freedom to include any number of blocks and filemarks within a set, and the ability to search to the end of it in one motion. Searching to setmarks automatically invokes the drive's fast-search facility, unless the setmark is in the buffer. The host does not

Tape Layout

need to know the number of blocks and filemarks contained in a set in order to position past it, before appending more data.

The use of the setmark is not restricted to marking the end of a set of data. The host is free to assign any meaning to this mark it wishes.

See "Managing Blocks and Separator Marks" in chapter SF1 for more information.

EOD Area

The EOD area specifies the point on the tape where the host stopped writing data. The host does not specifically command the writing of this section of the tape. It is the *drive's* responsibility to detect conditions which indicate that the host has stopped writing data, and generate the EOD area at this point.

EW and EOP/M

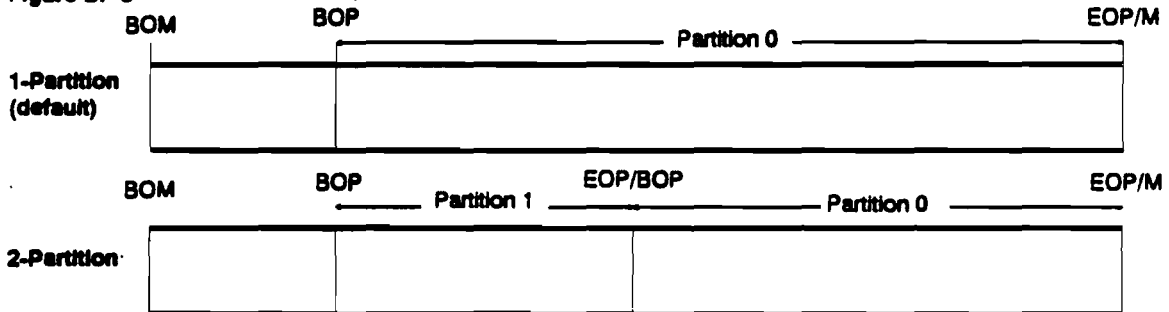
At the end of the tape are the Early Warning (EW) mark and End-of-Partition/Medium (EOP/M). EW is a fixed distance (approximately 500 mm) from EOP/M, and is generated automatically by the drive. When the drive detects the EW point while writing, it indicates to the host that it should stop writing data to the tape. After EOM, there are about 2 m of tape before the physical end of tape (EOT)—the point at which the leader tape is joined.

Tape Partitions

The DDS format provides the option of formatting the tape into a 1-Partition or 2-Partition structure. The host decides whether a blank tape is to be formatted into one or two partitions before writing any data. If no format command is sent by the host before writing to a blank tape, the tape will default to a 1-Partition structure.

If two partitions are created, the partition closest to the beginning of the tape is known as Partition 1, and the partition closest to the end of the tape is known as Partition 0. The size of Partition 1, in megabytes, is determined by the host during formatting (see "Making or Modifying Two Partitions" in chapter SF1). Each partition may be written and read independently. Figure DF-5 shows the available partition structures.

Figure DF-5



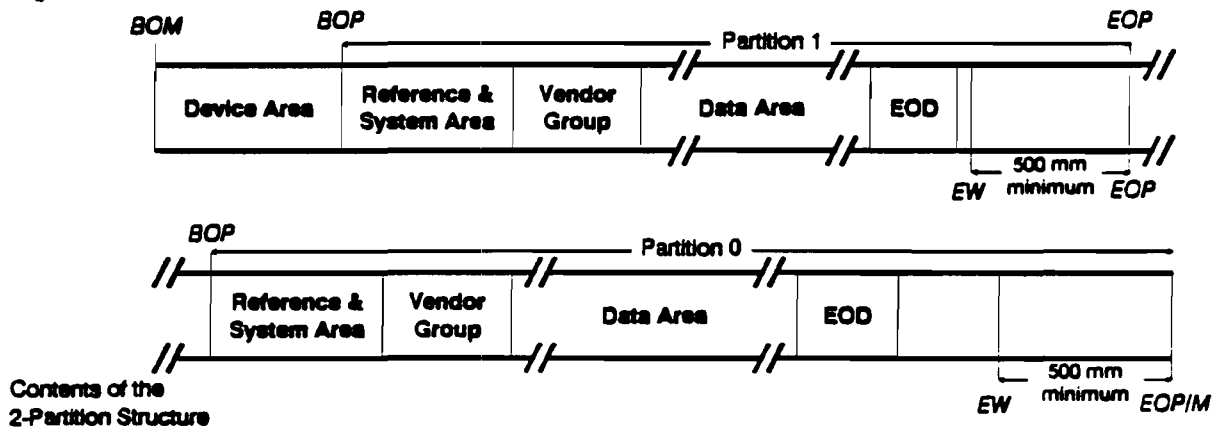
Available Partition Structures

**Note** In each of the two partitions, new data may be appended to existing data at any time. However, when data is written to a partition, any existing data which is beyond the current writing point in that partition will be lost.

## Two-Partition Structure

A 2-Partition tape has a single Device area, the same as a 1-Partition structure. Each partition has its own Reference and System area, Data area and EOD area. In addition, an EW mark and EOP are defined for Partition 1 and BOP is defined for Partition 0. See figure DF-6.

Figure DF-6



Two System areas are provided in order to maintain separate logs of usage and soft error occurrence for each partition. It is likely that the partitions will experience different degrees and types of use, and so show different error occurrences. The existence of two sets of logs allows the data to be retained separately for each partition; a single log might mask significant differences between the two partitions.

Each System area is preceded by a Reference area.

Separate Vendor Groups are written automatically by the drives that write each partition. They provide information about the drive, the time the partition was initialized, and the type of interface.

---

## Entities (DDS-DC)

After compression, data blocks are formed into collections, known as *entities*. Each entity also contains an 8-byte header (which is uncompressed), containing control information about the compressed data blocks within the entity. An entity contains one or more (up to 65,535) compressed data blocks.

In general, the existence of entities is completely transparent to the host system. However, in order for interchange to be possible between devices which support compression and those that do not, some host systems may implement software decompression. In this case, it is necessary for the host system to have a full knowledge of the structure of an entity.

The entity header is as follows:

Byte	Meaning
1	Bits 1-4 specify the header length (i.e. 8). Bits 5-8 are 0.
2	Reserved (0)
3	Registered algorithm identifier
4-6	Uncompressed length of every record in the entry
7-8	Number of records in the entity

Refer to the proposed ANSI standard "Helical-Scan Digital Computer Tape Cartridge, 3.81 mm (0.150 in), Extended Digital Data Storage Format (DDS-DC), for Information Interchange" for a full description of entities.


# Fast-Search, Indexing, Error Correction DF2

This chapter contains details of how fast-search, indexing and error correction techniques work in DDS-2 format drives. For details of how to implement these features, see chapter SF1.

## How Fast Search Works

The drive uses information encoded in the Sub-Code areas of each track (see figure DF-7) to perform fast-searching.

Figure DF-7




Location of Sub-Code Areas

The information in the Sub-Code areas includes a group count, setmark count, filemark count and block count. Fast-searching is done in response to a SPACE command from the host to find a specified setmark, filemark or block.

When fast-searching, the tape head follows a different path to that in the normal read mode, because both the head rotation rate and tape speed are changed. (See figure DF-8).

Figure DF-8



Path of head during Fast Search

To allow the drive's electronics to lock on and acquire the positional information required for fast-search, the information in the Sub-Code areas is repeated for each frame in a group.

Once the target group has been found, the drive uses the information in the group index to find out where specific setmarks, filemarks or blocks are located within the group.

Fast-searching is invoked by an internal algorithm making decisions about what is in the buffer, and depends on whether the command is a SPACE forward or reverse of a number of blocks, filemarks or setmarks.

# How Indexing Works

The method of indexing in DDS-2 used in a group allows the format to map variable-length blocks into fixed group sizes.

The index contains two special blocks:

1 *The Group Information Table*. This has a fixed size of 32 bytes, containing the following information:

- The Group Number
- The Block Count
- The Filemark and Setmark Counts
- The size of the Block Access Table

2 *The Block Access Table*. This is of variable length, and contains the following information:

- The start and end of entities and blocks in the group
- The position of filemarks and setmarks in the group


The minimum supported length for an entity or block is 1 byte

The maximum supported length for an entity or block is  $2^{24} - 1$  bytes

This limit is imposed by the 3 bytes allocated in the Block Access table for block length.

Figure DF-9 shows an example of Indexing, and table DF-1 shows the Block Access entries which correspond to the example.

Figure DF-9



Example of Indexing

Table DF-1. Block Access Entries

	Flag Byte	Count (dec)	Remarks
Access Entry 1	01100011	10240	10 Kbyte Block
	00000111	0	Filemark
	01100011	20480	20 Kbyte Block
	00000011	0	Filemark
	01110011	65536	64 Kbyte Entity
	00000111	0	Filemark
	01100011	10240	10 Kbyte Block
	00000111	0	Filemark
	00000111	0	Filemark
	00000111	1	Setmark
Access Entry 11	10000000	20136	Skip Entry

The skip entry is computed as the difference between group capacity and total count. In the above example:

$$\text{skip entry count} = 126632 - (10240 + 20480 + 65536 + 10240) = 20136.$$

The figure of 20136 includes the 32 bytes of the Group Information Table and the 44 bytes (11 × 4) of the Block Access Table entries.

## Block Access Table

The Block Access Table describes the contents of the group and contains an entry for each block, entity, filemark and setmark. Partial entities and blocks contained in the group also have an entry. The size of the table varies according to the contents of the group. The access entries are in reverse order, with the latest entry first and the first entry last, immediately before the Group Information Table.

## Access Entries

Each access table entry is 4 bytes, consisting of a Flag byte followed by a 3-byte Count field:

Bit:	7	6	5	4	3	2	1	0	
0	Bit 7	Bit 6	Bit 5	Bit 4	After EW	Bit 2	Bit 1	Bit 0	← Flag Byte
1	(MSW)			Count					← Flag Byte
2									
3									

Flag Byte	Description	Start Block	End Block	Count (bytes)
0000x001	Total Count of Unprocessed Records if last BAT entry 0110x000: if last 2 BAT entries 0110x000, 1000x000:	previous <previous	current previous	in the whole record in the whole record
0000x111	Separator Mark			0 = Filemark, 1 = Setmark
0001x001	Total Count of Entity			in whole previous entity
0100x000	Middle Part of Unprocessed Record	previous	next	in part of record in current block
0100x010	First Part of Unprocessed Record	current	next	in part of record in current block
0101x000	Middle Part of Entity	previous	next	in part of record in current block
0101x010	First Part of Entity	current	next	in part of entity in current block
0110x000	Entire Unprocessed Record	current	current	in whole record
0110x011	Last Part of Entity	previous	current	in part of entity in current block
0111x000	Last Part of Unprocessed Record	previous	current	in part of record in current block
0111x011	Entire Entity	current	current	in whole entity
1000x000	Skip Entry			remaining in the group

After EW After Early Warning when the block was written.

- 1 Allows the drive to return EW status to the host during reading at the same point at which it was detected during writing.

## How Indexing Works

Both Block Begin and Block End are set for filemarks and setmarks. However, End of Data-Block is only set for the end of a block or entity of data, not for marks.

If Block End is set and Mark is not set, then the Count field contains the number of bytes in the block or entity.

If Block End is set but neither Mark nor Transfer Data are set, then the Count field contains the total number of bytes for a block or entity which spans one or more groups. The count must equal the sum of the Count fields of the constituent elements of the block or entity, specifically the first part, any middle parts, and the last part of the block or entity.

The maximum supported length for a block or entity is  $2^{24} - 1$  bytes.

---

**Note** More detailed information on format is available in the DDS Format Description, available from Hewlett-Packard Computer Peripherals, Bristol, UK.

---

The maximum number of blocks, filemarks and setmarks that can be written is governed by the byte allocations in the Group Information Table.

These are as follows:

Blocks	$2^{32} - 1$	
Filemarks	$2^{32} - 1$	
Setmarks	$2^{16} - 1$	
Number of Groups	$2^{16} - 1$	(excluding the group designated as the Vendor Group, which is written automatically by the drive)

### Example of Indexing

The following calculation shows the effects of indexing on tape capacity when writing fixed-length 512-byte blocks. The example assumes no data compression. Capacity increases in direct proportion to compression.

	126632 bytes	Capacity of a group
-	32 bytes	Group Information Table (fixed overhead)
-	4 bytes	Skip Entry (fixed overhead)
	<hr/> 126596 bytes	Available data and index space

$$\begin{aligned} \text{Capacity} &= \frac{\text{available data and index space}}{\text{length of data block} + \text{block entry}} = \frac{126596}{512 + 4} \\ &= 245 \text{ data blocks} + 176 \text{ bytes extra} \end{aligned}$$

The 176 extra bytes for a partial block must include another index entry (4 bytes) leaving 172 bytes for the partial block data.

$$\begin{aligned} \text{Data size} &= 245 \text{ blocks} + 1 \text{ partial block} \\ &= (245 \times 512) + 172 = 125612 \text{ bytes/group} \\ \text{Index size} &= \text{Group Information Table} + \text{Skip Entry} \\ &\quad + \text{Index Entries for 245 blocks and 1 Partial block} \\ &= 32 + 4 + 4 + (245 \times 4) = 1020 \text{ bytes} \\ \text{Index overhead} &= \frac{1020}{126632} \times 100 = 0.8055\% \\ \text{Tape capacity} &= 2.6 \text{ gigabytes (nominal)} - 0.8055\% \text{ of } 2.6 \\ &= 2.579 \text{ gigabytes} \end{aligned}$$



## How Error Correction Works in DDS-2

This section reviews the ability of the Audio DAT format and the DDS-2 format to correct errors. Audio DAT has two levels of ECC—C1 and C2. DDS-2 format uses the same levels as Audio DAT format but adds extra error-correction techniques: C3 ECC, N-group writing, read-after-write, data randomizer and checksums.

### Audio Format C1 and C2 ECC Performance

C1 ECC is (32,28) Reed Solomon code and can detect and correct errors in any two symbols, or it can correct four symbols where the error location is known. A symbol is basically equivalent to one byte. C2 ECC is (32,26) Reed Solomon code and can correct errors up to three symbols long, or six symbols when the error location is known.

To minimize the probability of faulty detection of an error, the C2 code decoder typically corrects two symbols and six symbols where the error location is known, using C1 error condition flags. The research papers listed in the Bibliography review the performance. Their conclusions are summarized as follows:

### Random Error Performance


In the following,  $P_s$  = Symbol error rate.

Probability that a symbol error is uncorrectable:  $P_{un} = 3.8 \times 10^{23} P_s^{17}$

Probability that a symbol error is undetected:  $P_{mis} = 8.8 \times 10^{18} P_s^{17}$

This performance is shown in figure DF-10.

Figure DF-10



How Error Correction Works


Burst Error Performance

Maximum correctable burst error length = 792 symbols (6336 bits)

Since the linear recording density of DAT is 61,000 bits/in., this is equal to 2.64 mm along a helical track.

The horizontal correctable error width is  $2.64 \times \sin 6.25^\circ \approx 0.3$  mm, and the vertical correctable error width is  $2.64 \times \cos 6.25^\circ \approx 2.6$  mm (see figure DF-11).

Figure DF-11



DDS-Format C3 ECC Performance

In the DDS-2 format, C3 code can be added to improve data integrity. C3 ECC is (46,44) Reed Solomon code and can detect and correct errors as follows:

Correctable	Detectable but Uncorrectable	Misdetection or Miscorrection
1 error	1 error + 1 erasure error	$\geq 2$ errors
$\leq 2$ erasure errors	$\geq 3$ erasure errors	$\geq 2$ errors + 1 erasure error $\geq 1$ error + $\geq 2$ erasure errors

*error*: an uncorrected and misdetected, or miscorrected symbol error after C1/C2 correction. Refer to the curve "Misdetection" in figure DF-10.

*erasure error*: a symbol error detected but uncorrected by C1/C2 correction. Refer to the curve "Uncorrection" in figure DF-10.

Each C3 codeword symbol is interleaved and corrected by the Audio DAT C1/C2 code. As a result, the C3 code symbol error rate is the error rate after C1/C2 correction.

C3 code performance can be estimated as follows:

Random Error Performance

Probability of a symbol being uncorrectable by C3 ECC after C1/C2 correction:

$$P_{unC3} \approx 46C_3 P_{un} \approx 1.5 \times 10^4 P_{un} \quad (\text{See figure DF-12})$$

where  $P_{un}$  is the error rate after C1/C2 correction.

Figure C-3 shows that random error performance is more than adequate. The shaded area represents the actual bit error rates seen in media testing.

### Burst Error Performance

The most dangerous errors are burst errors caused by horizontal, vertical and helical error patterns. C3 can correct any two tracks in a group (see figure DF-13).

Figure DF-12





Figure DF-13



The code copes with the errors as follows:

**Horizontal error pattern**

C1/C2 code can correct up to 0.3 mm width horizontal error pattern. If the error exceeds this limit, all the tracks become uncorrectable, and so no type of C3 ECC can help.

**Vertical error pattern**

C1/C2 code can correct up to 2.6 mm width vertical error pattern. Outside this limit, about 280 tracks become uncorrectable simultaneously, assuming circular defects. Correcting an error like this using a track-based correction scheme would require a C3 capable of correcting more than 280 tracks.

**Helical error pattern**

C3 is well suited to correct helical errors (soft errors caused by head clog). One way of removing the error is through 'retry', but this affects the data transfer rate.

How Error Correction Works


Using C3 code:

- The drive can continue streaming.
- Hard errors which may occur during writing without read-after-write can be detected and corrected.

The chance of a head clog occurring depends very much on the condition in which the drive is maintained.

As a worst case, consider figure DF-14. 84 head clogs were observed during 3000 hours reading—in other words,  $1.6 \times 10^{13}$  bits read without maintenance. The probability of needing a retry is 23%, that is 19 cases during the last 1400 hours reading. This gives a rate of once in each 74 hours use, and this is quite acceptable in a worst case where there is no maintenance. With proper maintenance, the retry rate will be much lower.

Figure DF-14



Head Clog Data

For the data illustrated in this graph, a new tape was used, written to once and then put through 1500 read passes (FWD for 2 hours, then REW). A Sony DTC-1000 tape drive was used with a 4-head drum. It was cleaned before the first pass; after that there was no cleaning and no maintenance.

Because of the nature of head clogs, there is no limit to their length. C3 can only reduce the occurrence of retries.

### N-Group Writing Performance

N-group writing is an additional technique to improve data integrity with or without C3 code. Each group is duplicated N times, and the probability that each symbol of the C3 code becomes uncorrectable is improved as follows:

$$P_{un:N-group} = P_{un}^N$$

Random error performance for 2-group writing is shown in figure DF-12 as well.




This technique is beneficial, especially when writing without read-after-write, and can recover up to  $23(N-1)$  tracks head clog per head, and up to  $46(N-1)$  tracks tape damage. Thus 2-group writing can recover up to 23 tracks of head clog and 46 tracks of tape damage, 3-group up to 46 and 92 tracks respectively, 4-group up to 69 and 196, and so on,

up to 8-group, which could recover up to 161 tracks head clog per head and 322 tracks of tape damage.

### Read-After-Write

This is the most popular and proven method of improving data integrity. In the DDS-2 format, frames are rewritten because this is more efficient than rewriting groups (see figure DF-15).

Figure DF-15



A bad frame can be rewritten after up to seven other frames have been written. This allows a variety of read-after-write sequences with existing and future LSI chips, and using mechanisms with a drum diameter of 30 mm and a 90° wrap angle or a drum diameter of 15 mm and a 180° wrap angle.

Attempts can also be made to rewrite each frame up to 255 times to allow the drive to skip damaged areas equivalent to several tracks.

## Data Randomizer

The error rates between random data and worst-pattern data differ by a factor of 10. By using a data randomizer, it is possible to reduce the worst-case error rate. The randomizer also produces data with a spread of transitions which have a more consistent RF envelope, allowing accurate RF level detection as a read-after-write criterion.

## Checksum

If a head clog occurs during writing, there is a chance that the previously recorded data is not overwritten but remains intact. This is called *drop in*. It is important to check for this type of error occurring during both writes and reads, because it will not be detected by the track-based C1 and C2 codes.

The structure of the C3 ECC is such that it will correct any two tracks if the location of the erroneous tracks is known, but only one track if the track location is unknown. By recording the track checksum in four subcode areas, the location of a drop-in is sure to be detected using the track checksum. This allows us to provide a full two tracks-worth of ECC with the C3 algorithm.

---

---

**Data Compression Technique**

---

## Data Compression Technique Contents

DC1	Data Compression	
	Redundancy	1
	Choosing a Method of Data Compression	1
	Data Integrity	3
	Advantages of Data Compression	3
	The Effect on Fast-Search	3
	Uncompressed Data	4
DC2	The DCLZ Algorithm	
	The Compressed Output	5
	The Dictionary	6
	How DCLZ Compresses Data	6
	How DCLZ Decompresses Data	9



This chapter provides a background to data compression available in the HP C1533A, and explains why Hewlett-Packard developed the DCLZ algorithm for implementing it. The way the algorithm works is described in chapter DC2.

---

## Redundancy

Almost all data contains redundancy in the form of repetition. Redundancy is vital in human communication because of our imperfect attention span and the generally noisy communication environments. With computer data transactions, redundancy of this kind is merely wasteful. Lossless data compression is a means of stripping it out in such a way that it is possible to reconstruct the data exactly as it was originally, in other words, to put the redundancy back.

Redundancy occurs because most data exhibits pattern to some degree, in other words, repeated groups of symbols. If the data is truly random, there is no pattern. There is then no redundancy. In compressing data, the ideal is to achieve compressed data that has no redundancy, so that it possesses the qualities of truly random data.

---

## Choosing a Method of Data Compression

The effectiveness of any method of data compression depends on several factors. In developing a method for use in its drives, Hewlett-Packard had to find a balance between the following:

- What degree of compression was needed?

The measure used was the Compression Ratio, which is the ratio of the amount of uncompressed data to the amount of compressed data. Obviously the best possible ratio was desirable, but only after the following factors had been considered.

- How much time could the compression and decompression processes take?

It was decided that data compression and decompression should not impede the rate of data transfer, so that compressing and storing data should be at least as fast as storing the data without compressing it.

- Could any assumptions be made about the data?

If data is known to be of a particular kind, for example numerical, a particular algorithm may out-perform all other algorithms. Hewlett-Packard felt that its chosen method should be able to cope effectively with any kind of data, even if that meant it was not optimal for a particular type of data. This meant that the method had to be adaptive—able to react to the peculiar patterns of data as it encountered it.

Choosing a Method of Data Compression

- Could any sacrifice be made as to data integrity?

Clearly the answer had to be no. If anything, the compression technique should improve data integrity.

In summary, the chosen data compression algorithm had to satisfy the following criteria:

- It must not impede data transactions with the host.
- It must be adaptive, reacting to any type of data.
- It must not affect data integrity.
- Given the above, it should compress the data as much as possible.

In order that the method would stand a reasonable chance of becoming an industry standard, HP started with the popular Lempel-Ziv (LZ) algorithms and worked on improving them.

LZ algorithms are basically of two types:

- LZ1 uses a 'sliding window' history of the data stream to be compressed. In this, as each byte is processed, a fixed-size window around the byte is searched to find matching patterns. In this way, the compression reacts extremely sensitively to changing data types.
- LZ2 and LZW algorithms compile a dictionary as they process the data and look up matching patterns in this dictionary. The Compression Ratio can also be monitored, and if it becomes poor because the data type has changed, a new dictionary can be started. Since the structure of the dictionary is amenable to efficient searching, the algorithms can provide a faster rate of compression than LZ1, while still adapting to changing data types.

Hewlett-Packard's algorithm, DCLZ, is based on LZ2/LZW. It is implemented on chip in the HP C1533A. The implementation includes monitoring the Compression Ratio, so that the algorithm can react quickly to different types of data. DCLZ allows the drive to achieve better transfer rates than drives without data compression.

Hardware data compression such as this, where the compression/decompression algorithm is built into the tape drive, means that the process is completely transparent to the host computer; the host is unaware that it is happening. It is also possible to implement data compression through software on the host, but this slows the host's transfer rate because it is having to perform compression computation on top of its usual tasks. It also hampers interchange of data because the same software must be present on any host wishing to retrieve the data.

The DDS-2 format includes data compression. It is described in more detail in the DDS Format section of this manual.

---

## Data Integrity

Most errors occur during the reading or writing of data from or to the media, not in the digital electronics section of the drive. In other words, the danger areas are those in which physical activity is involved and where environmental conditions can play a part. Data compression takes place purely electronically, and reduces the amount of physical activity that is needed to get the data onto tape. Consequently, the compression decreases the error rate in real terms. For example, if the error rate without compression is 1 in  $10^{15}$ , and the compression ratio is 3:1, the error rate with compression becomes 1 in  $3 \times 10^{15}$ —three times better. If an error does occur however, more data will be lost, because more data is written in the same area of tape.

The data compression algorithm itself is lossless; that is, it guarantees that what was compressed can be decompressed without error. Equally, this means that if an error exists in the data before it is compressed, the error will still be there after decompression.

---

## Advantages of Data Compression

- Data compression allows the drive to store more data on the same length of tape.
- It allows the drive to match the performance of higher transfer rate systems more closely. If the drive is slower than the incoming host data stream, compression can aid the performance.

Theoretically, the increase in transfer rate is the same as the Compression ratio, but in practice there is a dependency on other elements, such as disk access times, the file system used, the file size, and system performance.

---

## The Effect on Fast-Search

Blocks of compressed data are stored on tape in collections called *entities* (see the section “Entities” in chapter DF1). An entity may span several groups on the tape.

In looking for a particular block on a data-compressed tape, the drive can fast-search to the beginning of the entity in which the block is. It does this in the same way as a drive without data compression, by reading the sub-code areas. However, the search will be quicker than with uncompressed data, because compressed data is more compact and so there is less distance to travel.

Having found the entity, the drive moves to the nearest access point before the block and starts decompressing. It must start there, otherwise it cannot properly reconstruct the compression dictionary. If the block is some way after the access point, then this will take time.

---

## Uncompressed Data

The DDS-2 format allows the drive to cope with mixtures of compressed and uncompressed data. This means that group headers and labels can be written uncompressed, so that any DDS-format drive can read them. It also means that groups can be formed of mixtures of filemarks, setmarks, uncompressed blocks, and entities of compressed data.

The drive can read any DDS-1, DDS-DC or DDS-2 format tape regardless of whether data compression is switched on or off, and automatically decompresses any DCLZ data regardless of whether data compression is switched on or off. Decompression is performed in a way which is transparent to the host system.

If the drive encounters data compressed using a different algorithm, it will not usually be able to decompress it. It can, however, pass the data to the host for decompression in software, if the host supports this. The same is true when non-compressing DDS-format drives such as the HP 35470A encounter compressed data. On meeting a compressed entity, such a drive will pass it to the host for decompression.

The Hewlett-Packard implementation (DCLZ) of the Lempel-Ziv algorithm provides a means of compressing a stream of data without making assumptions as to what type of data it is. Basically it assigns a codeword (a numeric value) to each pair or string of characters it meets which it has not met before. As it compresses, it compiles a dictionary which matches the codewords to the strings they represent. Any later occurrence of the string is replaced by the same codeword.

---

## The Compressed Output

The output stream consists of a stream of codewords, each of which is a number in the range 0–4095. These codewords fall into three types—Control Flags, Encoded Bytes, and Dictionary Codes:

---

<b>Codewords 0–7</b>	<b>Control Flags:</b> <ol style="list-style-type: none"><li>0 <i>Dictionary Frozen.</i> Used when the compression algorithm has almost filled its dictionary and is taking too long to find free space for further entries. It tells the decompression algorithm that it need not create any new dictionary entries.</li><li>1 <i>Dictionary Reset.</i> Clears all the Dictionary Codes so that the algorithm can start building a new dictionary. This codeword would normally be sent at the start of an entity, but the algorithm may also use it at other times if it seems that compression is not good enough because the current dictionary no longer reflects the redundancy characteristics of the data.</li><li>2 <i>Increment Codeword Size.</i> Says that subsequent codewords sent in the output stream will be one bit longer than the current codeword size. After a Dictionary Reset, output codewords are 9-bit until this size is too small to contain the required Dictionary Code. At that point, codeword 2 is sent, and the codeword length immediately rises to 10-bit. As further codewords are generated, the codeword length rises to 11- and then 12-bit as required. This ensures that the codeword length is never more than the minimum necessary.</li><li>3 <i>End-of-Block.</i> Indicates that the next codeword represents the last character (if the codeword is an Encoded Byte) or character string (if it is a Dictionary Code) of the block.</li></ol> <p>4–7 Not used.</p>
<b>Codewords 8–253</b>	<i>Encoded Bytes.</i> These contain the values 0–255, representing single bytes in the input stream. They might be, for example, the ASCII values of characters in the input stream, so that the codeword for character “a” would be (ASCII value of “a” + 8) = 97 + 8 = 105.
<b>Codewords 254–4095</b>	<i>Dictionary Codes,</i> built up as data is processed. These contain a character, plus a pointer to another codeword. This leads to a linked chain of codewords which produces a string of characters. In effect, then, a Dictionary Code represents a string of characters.

---

## The Dictionary

The dictionary is formed during compression, and stored temporarily in RAM in the drive. It is never explicitly written to tape, because the dictionary can be reconstructed from the codewords during decompression.

When a Dictionary Code is formed, it contains a pointer to an earlier Dictionary Code or to an Encoded Byte, plus the byte value of a new character. In this way it starts a linked chain which builds up a string of characters. The new Dictionary Code is the codeword for that string.

Each dictionary entry is 23 bits long, as follows:

BR		
22-20	19-8	7-0
Three condition flags	Codeword representing the entry, or pointing to another Dictionary Code or Encoded Byte	Byte Value of the entry

---

## How DCLZ Compresses Data

When compressing data, the DCLZ algorithm is constantly trying to match the incoming data with strings of characters it has met before, and which it has stored as dictionary entries.

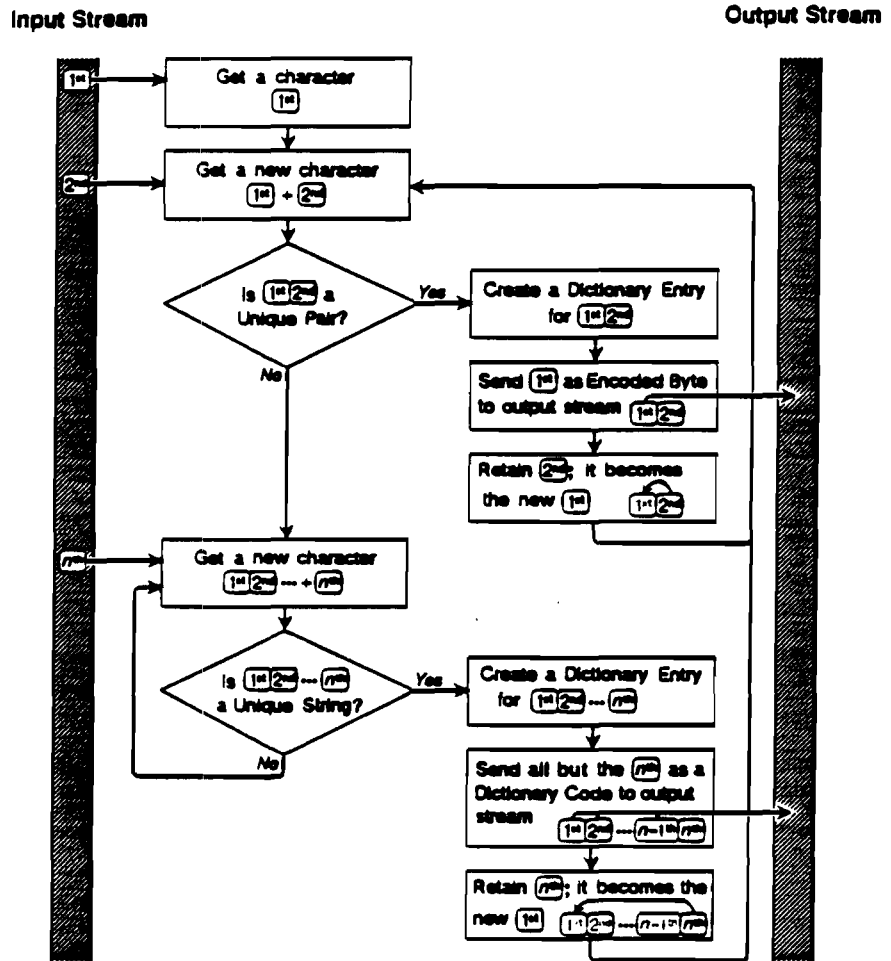
In the following explanation, a *Unique Pair* is a pair of characters which the algorithm has not met before, and which consequently has no entry in the dictionary.

Similarly, a *Unique String* is a string of more than two characters for which no dictionary entry exists.

- 1 The algorithm fetches characters one by one from the input data stream until it has either a *Unique Pair* or a *Unique String*.
- 2 A new dictionary entry is created for the pair or the string.
- 3 All but the last character of the pair or string is sent to the output stream as either an Encoded Byte or a Dictionary Code. The last character then becomes the starting point for a new search for a *Unique Pair* or *Unique String*. The algorithm, in other words, goes back to step 1.

This basic procedure is shown as a flow-diagram in figure DC-1.

Figure DC-1



Flowchart of the  
DCLZ compression  
process

It follows from this process, that if the algorithm found a Unique Pair, a single character is sent to the output stream. This will be a codeword in the range 8-263, an Encoded Byte. If a Unique String is found, the string minus its last character is sent to the output stream. This is in the form of a Dictionary Code, which provides the first step in a series of linked Dictionary Codes which will reproduce the string. The last in the series will be an Encoded Byte.

The DCLZ Algorithm  
How DCLZ Compresses Data

To provide an example which is reasonably short, suppose the data consists of the short string "RINTINTIN". The DCLZ algorithm would compress it as follows:

Input Stream	Current String	Unique Pair?	Unique String?	Dictionary Entry	Output	Meaning of Output
					1	Reset the dictionary
R	R	No	No			Clear and initialize the dictionary. Fetch the first character from the input stream.
I	RI	Yes	No	264 RI	90	Encoded Bytes for "R" "RI" is a unique pair, so create a dictionary entry for it. Output "R" as an encoded byte. Retain "I" as the current string, and fetch the next character from the input stream.
N	IN	Yes	No	265 IN	81	Encoded Bytes for "I" "IN" is a unique pair, so create a dictionary entry for it. Output "I" as an encoded byte. Retain "N" as the current string, and fetch the next character from the input stream.
T	NT	Yes	No	266 NT	86	Encoded Bytes for "N" "NT" is a unique pair, so create a dictionary entry for it. Output "N" as an encoded byte. Retain "T" as the current string, and fetch the next character from the input stream.
I	TI	Yes	No	267 TI	92	Encoded Bytes for "T" "TI" is a unique pair, so create a dictionary entry for it. Output "T" as an encoded byte. Retain "I" as the current string, and fetch the next character from the input stream.
N	IN	No	No			"IN" is not a unique pair since a dictionary entry already exists for it. So simply fetch the next character from the input stream.
T	INT	No	Yes	268 INT	265	Dictionary Code for "IN" "INT" is a unique string, so create a dictionary entry for it. Output "IN" as a Dictionary Code value. Retain the final "T" as the current string, and fetch the next character from the input stream.
I	TI	No	No			"TI" is not a unique pair since a dictionary entry already exists for it. So simply fetch the next character from the input stream.
N	TIN	No	Yes	269 TIN	267	Dictionary Code for "TI" "TIN" is a unique string, so create a dictionary entry for it. Output "TI" as a Dictionary Code value. Retain the final "N" as the current string, and observe that it is the last character in the input stream.
	N				3	End of block Output the end-of-block code, 3, to say that the next character is the last in the block.
	N				88	Encoded byte for "N"

Finally output the "N" as an encoded byte.

The output stream is then 1, 90, 81, 86, 92, 265, 267, 3, 86, each sent as 9 bits. Ignoring the two control characters, which overbalance this artificially short example, this is  $7 \times 9 = 63$  bits compared with the  $9 \times 8 = 72$  bits of the original—a compression ratio of 1.14:1. Typical compression ratios for substantial quantities of data are in the range 2:1 to 4:1.

The dictionary entries are shown as strings in the example above for clarity. In fact each entry consists of the codeword pointing to the first part of the string, plus the byte value for the last character. Thus entry 268 would be (265, T). Entry 265 would be (81, N), and 81 is the Encoded Byte for "I". In this way, the entries form a linked chain leading to the string "INT". The length of each dictionary entry however remains constant and minimal.



## How DCLZ Decompresses Data

The decompression algorithm starts with a reset dictionary containing only the Control Codes and Encoded Bytes. After that, it simply takes codewords from the input stream and looks them up in the dictionary. It then builds new Dictionary Codes combining the previously received codewords. In this way it reconstructs the dictionary created during compression, so that any codeword that is received is guaranteed to be in the dictionary.

The following example shows how the algorithm would decompress the sequence resulting from the compression example above. If codewords are simple Encoded Bytes, they are output straight away. If they are Dictionary Codes, the dictionary entries lead the algorithm through a series of bytes and codewords which point to other dictionary entries. The bytes are put on the stack until an Encoded Byte is encountered. Then the entire stack is output.

Input Stream	Previous Codeword	Codeword	Byte	Stack (LIFO)	Is Codeword an Encoded Byte?	Dictionary (Prev. CW, Byte)	Output
1	...	1	...	...	...		
Reset the dictionary.							
90	...	90	R	R	Yes	...	R
Look up the codeword (90) in the dictionary, and put the byte value (R) on the stack. Since the codeword is an Encoded Byte, output the stack and clear it. (Normally a dictionary entry would be created, but this is the first codeword after a dictionary reset.)							
81	90	81			Yes	264 (90, I)	I
Look up the codeword (81) in the dictionary, and put the byte value (I) on the stack. Since the codeword is an Encoded Byte, output the stack and clear it. Create a dictionary entry, using the Previous Codeword (90) and the Byte found in the dictionary (I).							
86	81	86	N	N	Yes	265 (81, N)	N
Look up the codeword (86) in the dictionary, and put the byte value (N) on the stack. Since the codeword is an Encoded Byte, output the stack and clear it. Create a dictionary entry, using the Previous Codeword (81) and the Byte found in the dictionary (N).							
92	86	92	T	T	Yes	266 (86, T)	T
Look up the codeword (92) in the dictionary, and put the byte value (T) on the stack. Since the codeword is an Encoded Byte, output the stack and clear it. Create a dictionary entry, using the Previous Codeword (86) and the Byte found in the dictionary (T).							
265	92	265	N	N	No		
Look up the codeword (265) in the dictionary. It is one of Dictionary Codes just created, and provides a byte value of "N" and a codeword of 81. Put "N" on the stack. Since 265 is not an Encoded Byte, nothing is output, and 81 now becomes the codeword.							
		81	I	NI	Yes	267 (92, I)	IN
Look up the codeword (81). It is an Encoded Byte, so add the byte value (I) to the stack. Output the stack (giving "IN" since the stack is LIFO) and clear it. Create a new dictionary entry, using the Previous Codeword (92) and the byte value (I).							

The DCLZ Algorithm

How DCLZ Decompresses Data

Input Stream	Previous Codeword	Codeword	Byte	Stack (LIFO)	Is Codeword an Encoded Byte?	Dictionary (Prev. CW, Byte)	Output
267	265	267			No		
Look up the codeword (267) in the dictionary. It is one of Dictionary Codes just created, and provides a byte value of "I" and a codeword of 92. Put "I" on the stack. Since 267 is not an Encoded Byte, nothing is output, and 92 now becomes the codeword.							
		92	T	IT	Yes	268 (265, T)	TI
Look up the codeword (92). It is an Encoded Byte, so add the byte value (T) to the stack. Create a new dictionary entry, using the Previous Codeword (265) and the byte value (T). Output the stack (giving "TI" since the stack is LIFO) and clear it.							
3	106	3	...	...	...		
The next entry is the last in the block.							
86	3	86	N	N	Yes		N
Look up the codeword (86) in the dictionary. It is an Encoded Byte giving a byte value of "N". Since this is the last codeword, simply output the byte value.							

In this way, the decompressed string "RINTINTIN" is recovered.

# Glossary

---

- access point** A point at the start of a sequence of compressed blocks, where the presentation of codewords to a decompression algorithm must start, even if the data required is not at the beginning of the sequence. Creation of a new dictionary will always begin at an access point, and an access point means the dictionary must be reset, although resets may occur at other times.
- AFC** Absolute Frame Count.
- algorithm** A rigorous set of rules for a procedure. In the context of data compression, the rules are for transforming the way data is represented.
- amble** A frame used to separate groups.
- ANSI** American National Standards Institute, which sets standards for, amongst other things, SCSI and the safety of electrical devices.
- area ID** Defines the area of the tape and specifies which type of frames are written.
- ATF** Automatic Track Following—a method of ensuring the head is in the center of the track being read. There are ATF areas at the beginning and end of each track, with the user data in between. The head is wider than a track and so reads the ATF areas of the two tracks on either side as well as that of the track it is supposed to be reading. By comparing the content and strength of the signals from the ATF areas, the servo system adjusts the drum and tape speeds so that the read-head passes directly over the center-line of the track being read.
- autoload** When a tape cartridge is inserted, a tape drive with autoload will automatically load it, without the host having to send a load command. If a drive does not have autoload, the drive will take no action until it receives a load command.
- BAT** Block Allocation Table
- bit error rate** 
$$\frac{\text{Number of errors}}{\text{Number of bits written or read}}$$
- block** A logical unit of information. Called “record” in the DDS-format specification.
- block count** The number of blocks written since the beginning of the tape. Filemarks and save-set marks are included in the Block Count.
- BOM** Beginning Of Media. The first point on the tape which can be accessed by the drive.
- BOP** Beginning Of Partition. The position at the beginning of the permissible recording region of a partition.
- buffered mode** A mode of data transfer in write operations which facilitates tape streaming. It is selected by setting the Buffered Mode Field to 1 or 2 in the SCSI Mode Select Parameter List header.
- burst error** A series of contiguous symbols on the tape which are incorrect.

- C1 ECC** C1 error correction code is (32,28,5) Reed-Solomon code with an interleave depth of two bytes. This enables it to correct up to two byte error or burst errors up to four bytes long. The code is stored in the same track as the data.
- C2 ECC** C2 error correction code is (32,26,7) Reed-Solomon code with an interleave depth of four blocks (1 data block = 288 data bits). This enables it to correct up to a three byte error, six byte erasure error, or 792 byte burst error. The code is stored on the same track as the data.
- C3 ECC** A third level of error correction code covered by the DDS format. C3 allows any two tracks in a group to be corrected, and is used only when a raw data error is too big to be corrected by C1 and C2. C3 code is stored in an extra frame at the end of the twenty-two frames of data in each group.
- checksum** The sum of a series of bytes written to the tape, which can be checked against the sum of the same series of bytes when the tape is read.
- codeword** A number between 0 and 4095 which is a reference to a dictionary entry for a sequence of one or more uncompressed data bytes. For example, the words "the theory of the pantheon" could be represented as "4 4ory of 4 pan4on" if it was understood that the codeword "4" represented "the". Alternatively, a codeword can contain other information which is used for the management of decompression and data retrieval, for example "1" means "reset the dictionary". A codeword can be from 9 to 12 bits long.
- compression** A procedure in which data is transformed by the removal of redundant information in order to reduce the number of bits required to represent the data. This is done by representing strings of bytes with codewords.
- compression ratio** A measure of how much compression has occurred, defined as the ratio of the amount of uncompressed data to the amount of compressed data into which it is transformed. The HPLZ algorithm can typically achieve a compression ratio of between 2:1 and 4:1 depending on the nature of the data.
- crosstalk** The condition in which the signals from one track on a tape interfere with the signals from an adjacent track.
- DAT** Digital Audio Tape
- data format ID** Specifies which data format is being used on the tape.
- decompression** A procedure in which codewords are transformed to regenerate the original representation of data.
- DCLZ** Data Compression Lempel-Ziv. A compression algorithm based on the Lempel-Ziv LZ2/LZW algorithms but with improved performance.
- device area** The first area on the tape used by the device for drum spin-up and testing.
- DDS** Digital Data Storage, a standard format originally developed by Hewlett-Packard and Sony for DAT used for data storage. The first generation standard is DDS-1 (or simply DDS), to which was added data compression to produce the DDS-DC standard. Further enhancements, notably narrower tracks and thinner tape, have led to DDS-2, the third generation, which can provide typically double the capacity of DDS-1.
- DDS-DC** An extension of the DDS format to add support for in-drive data compression and decompression. The support is intended for lossless algorithms such as those of the Lempel-Ziv family, but is not limited to them.
- dictionary** A map which relates codewords to the uncompressed data from which they are derived. It is created during the compression of data, but not stored

- explicitly on tape. During decompression, the directory is recreated from the codewords on tape.
- dictionary code** A dictionary entry in the codeword range 264–4095. Dictionary Codes are built up as compression occurs, as opposed to Encoded Bytes, which are non-adaptive. Each Dictionary Code contains a character, plus a codeword pointing to another Dictionary Code or to an Encoded Byte. The entry forms the start of a linked chain of entries which resolve into a string of characters. The chain ends when it reaches an Encoded Byte, which represents only a character.
- dropout** An area of tape where the signal level of the medium has fallen off to a level where data recovery is no longer possible.
- ECC** Error Correction Code. See C1 ECC, C2 ECC and C3 ECC.
- ECMA** European Computer Manufacturer's Association
- encoded byte** A dictionary entry in the range 8–263 which contains the byte value of a character and corresponds to that character, as opposed to a Dictionary Code, which also points to another dictionary entry.
- entity** A unit of recorded information containing a sequence of equal-length blocks which have then been compressed using the same algorithm. The entity can span groups on the tape. It contains at most one access point, which will be at the start of the first complete block in the entity.
- envelope** A waveform composed of the instantaneous peak values of an alternating signal which indicates the variation in peak amplitude of the signal.
- EOD** End Of Data
- EOM** End Of Media format. The last usable point on the tape.
- EOP** End Of Partition. The position at the end of the permissible recording region of a partition.
- EW** Early Warning. A physical mark or a device computed position on the tape which tells the drive that it is approaching EOP.
- fast searching** The process of reading just the sub-areas to locate an item on the tape at a speed of 1630 mm/s.
- FRU** Field Replaceable Unit, an assembly or group of components which is replaced in its entirety by Service Engineers when it contains a fault.
- filemark** A mark written by the host. It does not necessarily separate files. It is up to the host to assign a meaning to the mark.
- filemark count** The number of filemarks written since the beginning of the current partition up to and including the current group.
- frame** Two adjacent tracks, one positive azimuth and one negative azimuth.
- group** A fixed capacity set of frames written to or read from the tape, defined in the DDS format.
- group count** The number of user data groups that have been written following the Vendor Group, starting with one. The Vendor Group has a group count of zero.
- hard error** An uncorrectable data error. During writing, this is defined as being uncorrected after the RAW retry limit has been exceeded. During reading, a hard error is logged if a group is uncorrectable by C1, C2 or C3 ECC.

- head clog** Particles from the tape or from outside the drive adhere to the head gap on a read or write head and obstruct the reading or writing of data. The particles will often become dislodged again with continued use.
- index** Information stored at the end of a group which specifies the contents of the group. Every group except the Vendor Group contains one index.
- interleaving** The process of shuffling the order of data bytes before writing them to tape so that consecutive bytes are recorded as far away from each other as possible. This minimizes the impact of any burst error, so that C1 and C2 ECC have the maximum chance of recovering the data.
- lead-in area** The first section of the tape used for loading, BOP positioning, and tape usage logging.
- LIFO** Last in, first out. Data written to a LIFO buffer is retrieved in the reverse order from which it was written.
- load** The process in which the drive takes in an inserted cartridge, opens the lid, and goes online.
- logical frame number** A unique number for each unique frame of a group. An amble frame has a Logical Frame Number of zero. Any repeated frames will have the same Logical Frame Number.
- LUN** Logical Unit Number, by which a device is identified on the SCSI bus. The drive has a fixed LUN of 0.
- LZ** Lempel-Ziv. A family of compression algorithms, all of which remove redundancy from data by encoding the data through a dictionary which is implicit in the compressed data.
- Media Recognition System** A system developed by Hewlett-Packard so that drives can recognize tapes which meet the DDS specification. Qualified tapes carry stripes on the transparent leader tape at the beginning of the tape which the drive can identify. The recognition system is switchable on the HP C1533A drive.
- N-group writing** Sometimes called multiple group writing, N-group writing repeats each group of data so that there are N consecutive copies of each group on the tape. This is a simple way of improving data integrity, but speed and capacity are sacrificed in writing all data several times.
- noise** Any kind of unwanted magnetic or electric interference detected by the electronics.
- online** The drive is online when a tape is loaded. The host has access to all command operations, including those which access the tape, set configurations and run diagnostic tests.
- offline** The drive is offline if the tape is currently unloaded or not in the drive. The host has limited access, and cannot perform any commands which would cause tape motion. The host can, however, load a tape, if one is inserted, and can execute any diagnostic tests which do not require tape motion.
- partition** A part of a tape which can be treated as a complete and independent whole. A tape can have one or two partitions.
- QIC** Quarter-Inch Cartridge. An industry committee for standardization of Quarter-Inch Cartridge formats.

- randomizing** A recoding of data symbols before they are written to tape in order to provide a consistently high RF envelope level. A consistent RF envelope is one of the criteria for rewriting a frame on read-after-write.
- RAW** *see* Read-After-Write
- raw bit error rate** The probability of a bit being an error, without using any error correction techniques (*see* bit error rate).
- read-after-write** RAW improves data integrity by reading data immediately after it is written and writing the frame again if an error is found. The Audio DAT two-head drum is replaced by a four-head drum for this, with two read-only heads and two write-only heads. Frames are only re-written as necessary, so speed and capacity are affected minimally. RAW is included in the DDS format.
- record** *see* block. "Record" is the DDS term for "block".
- redundancy** Repetition within data, where a string of characters occurs with greater frequency than the norm. By replacing the string with a single symbol, the data can be compressed without losing any information.
- reserved** Not generally available for use with the drive. A reserved field should contain all zero bits.
- setmark** A special recorded element within a partition, which allows fast searching to a point on the tape without having to know the number of records or filemarks that precede this point.
- setmark count** The number of setmarks that have been written since the beginning of the current partition.
- SCSI** Small Computer System Interface
- sliding window** A technique of data compression, used in the LZ1 algorithm, where redundancy is sought within a window of data centered on the byte currently being processed. This means that the algorithm can respond very sensitively to changing patterns of repetition, but there is an overhead in the amount of time needed to perform the search over the window with each byte processed. DCLZ is not of this type.
- soft error** A soft error is a data error which can be corrected by a RAW rewrite during writing to tape, or by C1, C2 or C3 ECC, or a read-retry during reading.
- spacing** Spacing is moving along the tape over a specified number of blocks, filemarks or setmarks, or to EOD, in order to find the data you want quickly.
- stack** An area of memory used for the temporary storage of data during processing. It is designed on LIFO (last-in first-out) principles and is manipulated by instructions to push (add to the stack) or to pop (remove from the stack).
- sub-code area** Sub-code areas are at the ends of tracks on the tape, and are read during fast-search. They contain details of block count, setmark count and filemark count, so that fast-search can find a particular point on the tape quickly.
- system area** A section in the Lead-in Area used to store the tape usage information.

## Glossary

- tape log** The tape log contains details of the history of a tape, the total number of groups written, of RAW retries, of groups read, of C3 ECC retries, and of loads. The log is copied into RAM when the tape is loaded into the drive, updated as the tape is used, and loaded back into the System area on the tape when it is unloaded.
- tape mark** Same as filemark (filemark is the preferred term).
- unique pair** A pair of characters taken from the input stream during data compression for which there is no entry in the current dictionary.
- unique string** A string of three or more characters taken from the input stream during data compression for which there is no entry in the current dictionary. The nature of the compression algorithm is such that there must be a dictionary entry for the string formed by removing the last character of a unique string.
- vendor unique** The addition of commands to SCSI that are not included in the standard.
- write-behind** Also called "Immediate reporting during writes". In write-behind mode, a drive accepts data from a host and reports that it has been written to tape when the data is actually still in the drive's buffer waiting to be written. This has the advantage that the host can continue with other tasks, based on the assumption that the data has been written.
- write-behind condition** A condition where the drive's buffer contains data which it has reported as having been written, but in fact still remains to be written.
- write-behind queue** Data (host data, filemarks and setmarks) held in the drive's buffer which has yet to be written to tape, despite the drive having reported to the host that the data has been written.