background information

hp ultrium drives technical reference manual

generations 1 and 2 SCSI and FC drives

volume 6: background to ultrium drives

Part Number: **C7379–90900 Volume 6**

Edition 2, February 2003

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.

Revision History

Version	Date	Changes
Edition 1		All. HTML version.
Edition 2	Feb 2003	All.PDF version for Generation 1 and Generation 2 SCSI and FC drives.

This document is frequently revised and updated. To find out if there is a later version, please ask your HP OEM Representative.

The Purpose of this Manual

This is one of five volumes that document HP Ultrium drives. This volume provides background information for driver and application developers. The following products are covered. Capacities are when the drive is using data compression with a compression ratio of 2:1, where applicable:

- HP Ultrium Generation 1 Full-Height SCSI Internal Drive
- HP Ultrium Generation 1 Half-Height SCSI Internal Drive
- HP Ultrium Generation 2 Full-Height SCSI Internal Drive
- HP Ultrium Generation 2 Full-Height FC Internal Drive

Note Throughout this manual frequent reference is made to SCSI commands. For more information on SCSI commands for HP Ultrium drives see volume 3, The SCSI Interface, of the HP Ultrium Technical Reference Manual set. Ordering details are given below.

Related Documents

The following documents provide additional information:

Documents Specific to HP Ultrium Drives

- **Hardware Integration Guide**, volume 1 of the HP Ultrium Technical Reference Manual
- **The SCSI Interface**, volume 3 of the HP Ultrium Technical Reference Manual
- **Specifications**, volume 4 of the HP Ultrium Technical Reference Manual
- **HP Ultrium Configuration Guide**, volume 5 of the HP Ultrium Technical Reference Manual
- **Background to Ultrium Drives**, volume 6 of the HP Ultrium Technical Reference Manual

Please contact your HP supplier for copies.

- The features and benefits of HP Ultrium drives are discussed in the HP Ultrium Technology White Paper.

- For a general backgrounder on LTO technology and licensing, go to `http://www.lto-technology.com`.

Documentation Map

The following will help you locate information in the 6-volume Technical Reference Manual:

Drives—general

	SCSI Drives	FC Drives
Connectors	**1** HW Integration: *ch.* 7	**1** HW Integration: *ch.* 4
Controller architecture	**6** Background: *ch.* 4	
Front Panel LEDs	**1** HW Integration: *ch.* 6	**1** HW Integration: *ch.* 3
Mechanism and hardware	**6** Background: *ch.* 3	
Specifications	**4** Specs	

Installation and Configuration

	SCSI Drives	FC Drives
Connectors	**1** HW Integration: *ch.* 7	**1** HW Integration: *ch.* 4
Determining the configuration	**2** SW Integration: *ch.* 2	**2** SW Integration: *ch.* 2
External drives *(SCSI only)*	**1** HW Integration: *ch.* 5	n/a
In Libraries	**1** HW Integration: *ch.* 1	
In Servers *(SCSI only)*	**1** HW Integration: *ch.* 4	n/a
In Tape Arrays *(SCSI only)*	**1** HW Integration: *ch.* 3	n/a
Modes of Usage *(SCSI only)*	**1** HW Integration: *ch.* 8	n/a
Optimizing performance *(SCSI only)*	**1** HW Integration: *ch.* 8	n/a
	2 SW Integration: *ch.* 4	
UNIX configuration	**5** UNIX Config	

Operation

	SCSI Drives	FC Drives
External drives *(SCSI only)*	**1** HW Integration: *ch.* 5	n/a
In Libraries	**1** HW Integration: *ch.* 1	
In Servers *(SCSI only)*	**1** HW Integration: *ch.* 4	n/a
In Tape Arrays *(SCSI only)*	**1** HW Integration: *ch.* 3	n/a

Cartridges

	SCSI Drives	FC Drives
Cartridge Memory (LTO-CM)	**2** SW Integration: *ch.* 5 **6** HW Integration: *ch.* 5	
Cartridges	**1** HW Integration: *ch.* 9	**1** HW Integration: *ch.* 5
Features	**6** HW Integration: *ch.* 5	
Managing the use of cartridges	**2** SW Integration: *ch.* 1	
Use of cartridges	**2** SW Integration: *ch.* 3	

Interface

	SCSI Drives	FC Drives
SCSI Guide	**3** SCSI	
Commands	**3** SCSI: *ch.* 4	
Error codes	**1** HW Integration: *ch.* 10	**1** HW Integration: *ch.* 6
Implementation	**3** SCSI: *ch.* 1	
Interpreting sense data	**2** SW Integration: *ch.* 3	
Messages	**3** SCSI: *ch.* 2	
Mode pages —see the MODE SENSE command	**3** SCSI: *ch.* 4	
Pre-execution checks	**3** SCSI: *ch.* 3	
Responding to Sense Keys and ASC/Q	**2** SW Integration: *ch.* 6	
Sense Keys and ASC/Q —see REQUEST SENSE command	**3** SCSI: *ch.* 4	

Maintenance and Troubleshooting

	SCSI Drives	FC Drives
Cleaning	**2** SW Integration: *ch.* 5 **2** SW Integration: *ch.* 7	
External drives *(SCSI only)*	**1** HW Integration: *ch.* 5	n/a
In Libraries	**1** HW Integration: *ch.* 1	
In Servers *(SCSI only)*	**1** HW Integration: *ch.* 4	n/a
In Tape Arrays *(SCSI only)*	**1** HW Integration: *ch.* 3	n/a
Monitoring drive and tape condition	**2** SW Integration: *ch.* 7	
Software troubleshooting techniques	**2** SW Integration: *ch.* 1	

Dealing with Errors

	SCSI Drives	FC Drives
Error Codes	**1** HW Integration: *ch.* 10	**1** HW Integration: *ch.* 6
Handling errors	**2** SW Integration: *ch.* 5	
How error correction works	**6** Background: *ch.* 4	
Logs—see the LOG SENSE command	**3** SCSI: *ch.* 4	
Recovering from write and read errors	**2** SW Integration: *ch.* 7	
Software response to error correction	**2** SW Integration: *ch.* 3	
Software response to logs	**2** SW Integration: *ch.* 3	
TapeAlert log	**2** SW Integration: *ch.* 7	

Ultrium Features

	SCSI Drives	FC Drives
Adaptive Tape Speed (ATS)	**6** Background: *ch.* 1	
Autoload	1 HW Integration: *ch.* 2	
Automation Control Interface (ACI)	1 HW Integration: *ch.* 2 6 Background: *ch.* 1	
Cartridge Memory (LTO-CM)s	1 HW Integration: *ch.* 2 **2** SW Integration: *ch.* 5 **6** HW Integration: *ch.* 5	
Data Compression, how it works	**6** Background: *ch.* 5	
Data Compression, managing	**2** SW Integration: *ch.* 5	
Design principles	**6** Background: *ch.* 1	
OBDR and CD-ROM emulation	**6** Background: *ch.* 1 **2** SW Integration: *ch.* 7	
Performance optimization	**1** HW Integration: *ch.* 8 **2** SW Integration: *ch.* 1	n/a
Performance, factors affecting	**2** SW Integration: *ch.* 4	
Software design	**2** SW Integration: *ch.* 1	
Supporting Ultrium features	**2** SW Integration: *ch.* 5	
Ultrium Format	**6** Background: *ch.* 2	

General Documents and Standardization

- Small Computer System Interface (SCSI-1), ANSI X3.131-1986. This is the ANSI authorized standard for SCSI implementation, available through ANSI
- Enhanced Small Computer System Interface (SCSI-2), ANSI X3T9.2-1993 Rev. 10L, available through ANSI

Copies of General Documents can be obtained from:

ANSI	11 West 42nd Street New York, NY 10036-8002 USA	
ISO	CP 56 CH-1211 Geneva 20 Switzerland	
ECMA	114 Rue du Rhône CH-1204 Geneva Switzerland	*Tel:* +41 22 849 6000 *Web URL:* `http://www.ecma.ch`
Global Engineering Documents	2805 McGaw Irvine, CA 92714 USA	*Tel:* 800 854 7179 or 714 261 1455

contents

Ultrium Features

1

Design Principles

- Proven technologies:
  - The load/unload mechanism
  - Metal Particle (MP) media with a PEN substrate for stability over temperature variations
- The ability to handle 100% duty cycle
- A simple tape path and a robust head/media interface
- Rugged cartridge design, that focuses on eliminating reel lock and leader misalignment failures
- Efficient head cleaning:
  - There is an active, internal head cleaner.
  - Cleaning cartridges have LTO-CM.
- Efficient cooling:
  - The design requires low power (<30W average)
  - The air flow is under mechanism, across the PCA, not through the head/tape path area, avoiding contamination from air-borne particles.

Design for Usability

- Automatic softload mechanism:
  - No handles or buttons are needed to load a cartridge.
  - Cartridges are automatically ejected by 25 mm for easy retrieval.
- Easy configuration; SCAM provides automatic SCSI address selection.

- Clear user interface, with 4 front panel LEDs in the basic version. The focus is on clear indications of the status of the drive and when user attention is required.

Design for Performance

- Maximum throughput:
  - Native throughput, *Gen 1:* 15 MB/s; *Gen 2:* 30 MB/s
  - Hardware data compression
  - High bandwidth interfaces
    - FCAL = 2 GB/s
    - *Gen 1:* Ultra2 SCSI = up to 80 MB/s
      Gen 2: Ultra3 SCSI = up to 160 MB/s
- Data rate matching:
  - 8 channel, *Gen 1:* 5 to 15 MB/s, *Gen 2:* 10 to 30 MB/s
  - Features to enable the host to keep streaming:
    - Adaptive tape speed (ATS)
    - A reposition time of less than 2.5 seconds
    - Cache, *Gen 1:* 16 MB, *Gen 2:* 64 MB
- Fast time to data:
  - Fast load time (*Gen 1:* <25s, *Gen 2:* <19s), including soft load, tape threading, tuning and positioning the tape at BOT
  - The search speed (*Gen 1:* 4.1 m/s, *Gen 2:* 5.5–7 m/s) provides the following times:

Generation 1 drives:

tape	average time to data
200 GB*	71s

Generation 2 drives:

tape	average time to data
400 GB* (Gen 2)	52s
200 GB* (Gen 1)	52s

*capacities at 2:1 data compression

— The search speed (4.1 m/s) provides the following time:

tape	average time to data
200 GB*	71s

*capacity at 2:1 data compression

— The search speed (5.5–7 m/s) provides the following times:

tape	average time to data
400 GB* (Gen 2)	52s
200 GB* (Gen 1)	52s

*capacities at 2:1 data compression

- Other factors:
  - Cartridge memory provides instant access to tape directory, history, and so on
  - Fast unload time: *Gen 1:* less than 13 seconds, *Gen 2:* less than 19 seconds

Adaptive Tape Speed (ATS)

Adaptive tape speed (ATS) is a new feature specific to HP Ultrium-1 tape drives, sometimes called 'data-rate matching'.

HP Ultrium Generation 1 drives have a 16 MB data buffer. Generation 2 drives have a 64 MB data buffer. In order to 'stream' the drive successfully in either read or write mode, we have to maintain data in this buffer. By varying the tape speed we can help to maintain streaming by changing the rate at which the buffer empties. The benefit of this is that we can minimize or even eliminate repositions caused by the buffer emptying of data.

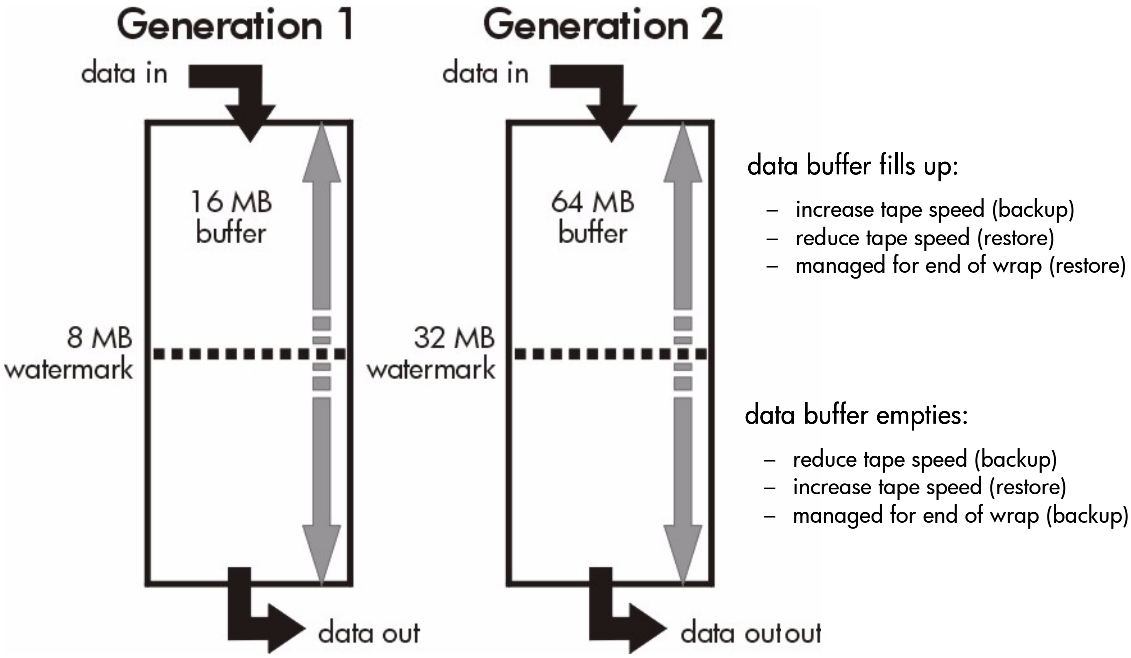

The illustration above shows how ATS works. The drive can vary its native transfer rate from its maximum speed (*Gen 1:* 15 MB/s, *Gen 2:* 30 MB/s) to a minimum of 6 MB/s (Gen 1) or 10 MB/s (Gen 2). By varying the speed of the tape, it attempts to maintain 8 MB (Gen 1) or 32 MB (Gen 2) of data in the data buffer, that is, half full. During a backup, as the buffer starts to fill above 8 or 32 MB, the tape speed is increased. This prevents us from having to disable host transfer if the data is coming in too quickly, helping us match the data-rate of the server.

If the buffer starts to empty below 8 or 32 MB, the tape speed is reduced. To achieve optimum performance, the data buffer is emptied as the end of wrap is approached. While the tape reverses and the head repositions the buffer is filled by the host. The buffer is then emptied to the 8 or 32 MB level when writing starts on the next wrap. When restoring data, the operation is reversed at the end of the wrap.

By reducing the number of repositions that occur during normal operation of the drive we can improve the overall reliability.

Automation Control Interface (ACI)

The Automation Control Interface (ACI) allows the activities of the drive to be coordinated within a library. The protocol has been designed so that it can be made into a standard feature of tape drives. It provides a rich and extensible functionality to allow automation manufactures to add value in their application of it.

The interface enables the following functionality for drives and libraries:

- It coordinates Load and Unload with a library.
- It provides information to the library.
- It enables the library to configure the tape drive.
- It enables firmware upgrades, in both directions.
- Though an extension, it allows a library to packetize SCSI commands for the tape drive to execute (SCSI Pass-Through mode).
- Though an extension, it allows the drive to act as host interface for a library (SCSI Surrogate mode).

HP Ultrium drives are equipped with an RS-422 serial port whose specification and access method are in the public domain. Primarily this allows for the attachment of robotic library controller devices.

Chapter 2, “Using Special Features in Libraries” in the **Hardware Integration Guide**, Volume 1 of the HP Ultrium OEM Technical Reference Manual has details of how the ACI can be used in tape libraries.

OBDR and CD-ROM Emulation

Note OBDR is only available on Generation 1 drives.

The One Button Disaster Recovery (OBDR) functionality in HP Ultrium drives enables them to emulate CD-ROM devices in specific circumstances (also known as being in “Disaster Recovery” mode). The drive can then act as a boot device for PCs that support booting off CD-ROM. In this way it is possible to recover from a “disaster” to the point at which normal applications will run with fully restored data, using only a single piece of media through a short, uncomplicated end-user process.

During CD-ROM mode, the Ready LED on the front panel alternates short and long flashes.

Entering CD-ROM Mode

A CD-ROM capable drive can be switched into CD-ROM mode by powering on with the Eject button held down. The drive then flashes the front panel LEDs in a warbling sequence. It will then respond to SCSI commands specific to CD-ROM mode and will identify itself as a CD-ROM device. For full details of the SCSI changes in CD-ROM mode, see “CD-ROM Emulation” in Chapter 1, “Interface Implementation” of **the SCSI Interface**, Volume 3 of the HP Ultrium Technical Reference Manual.

Writing is disabled. Use of the eject button at power-on is respecified to cause entry to CD-emulation mode.

When a tape is inserted (or present on entry) in CD-ROM emulation mode, it is assumed to contain an image of a CD offset 20 blocks into the tape. The drive reads the first 250 kilobytes of this CD image into buffer space reserved for CD-caching. It looks for a special message (`EL TORITO SPECIFICATION`) at the 8th byte of the 18th record of the CD image. If it cannot find this message, the drive ejects the tape automatically and waits (remaining in CD-emulation mode) for a properly written CD-image tape to be inserted.

Leaving CD-ROM Mode

The drive will remain in CD-emulation mode until one of the following occurs:

- A Mode Select command switches it back to tape drive mode using the CD-emulation mode page.
- A SCSI bus reset occurs following the reading of at least 100 blocks of CD-ROM data by a host.
- The user power-cycles the drive or resets it using the forced-eject mechanism.

Format

2

How Data Is Stored on Tape

This is an outline description of the Ultrium format as implemented in HP Ultrium tape drives.

The format was developed in partnership with the controller architecture, building on experience from DDS tape drives while considering potential formats for future generations of the HP Ultrium tape drive product family.

The following diagram showing the stage by stage transformations from SCSI records to the tape:

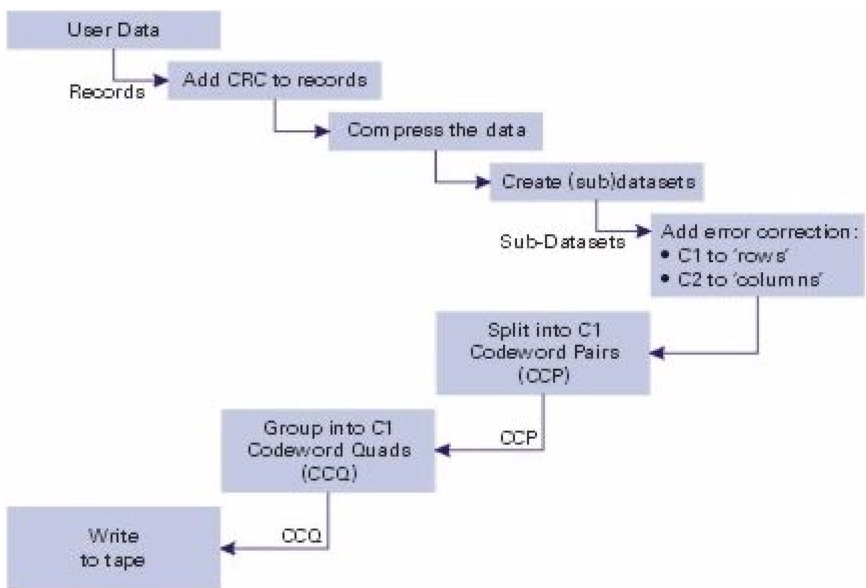

Compression

The data compression hardware in HP Ultrium drives can detect whether incoming data is already compressed and will not attempt to compress it again. The drive is able to switch silently and with great agility between compressing and non-compressing modes; in this way it optimizes both compression ratio performance and data rate. This mode of behavior is embedded in the Ultrium format.

The data is compressed using the LTO-DC compression format which is based on ALDC (licensed from Stac/IBM) with two enhancements:

- It uses two schemes to reduce expansion with uncompressible data. The scheme is selected automatically, based on the compressibility of the data currently being processed.
  - Scheme 1 is *compression*, using ALDC.
  - Scheme 2 is a *pass-through mode* that limits expansion to less than 0.5%.
- It uses embedded codewords.
  - These avoid the need for look-up table for records and filemarks.
  - They signal the fact that the processing engine has changed data compression schemes.
  - They indicate that the compression history buffer has been reset.

Error Detection and Correction

Once the data is compressed, error detection and correction information is added:

- Two levels of ECC are added to protect data:
  - 10^{-17} Uncorrectable
  - 10^{-27} Undetectable
- The format is designed to reduce susceptibility to scratches (both along and across the tape). Data can still be recovered from tape even if one head not working (on an 8-track device).
- Cyclic Redundancy Checks are added to records before compression for extra protection.
- Write Pass IDs are included to protect against old data showing through new data (drop-ins)

EEC Encoded Sub-Data Sets

The data is split up into 'sub-data sets' of 25,272 bytes and C1/C2 Error Correction Codes are added:

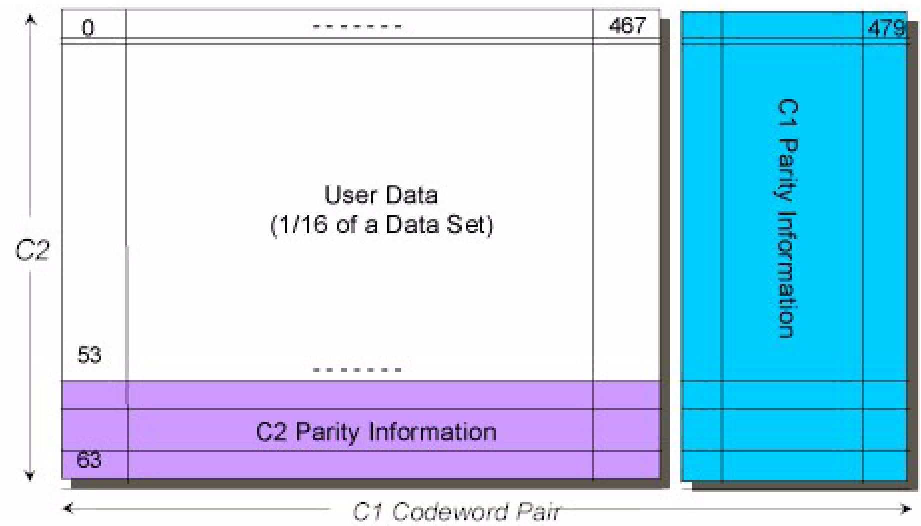

- A data set contains 404,352 user bytes; a sub-data set contains 25272 user bytes.
- C1 is a (240, 234, 7) code. Two codewords are interleaved on each row (480 bytes).
- C2 is a (64, 54, 11) code. Each column has one C2-codeword.
- 16 sub-data sets make a C2-protected data set.

The data is then split into C1 Codeword Pairs:

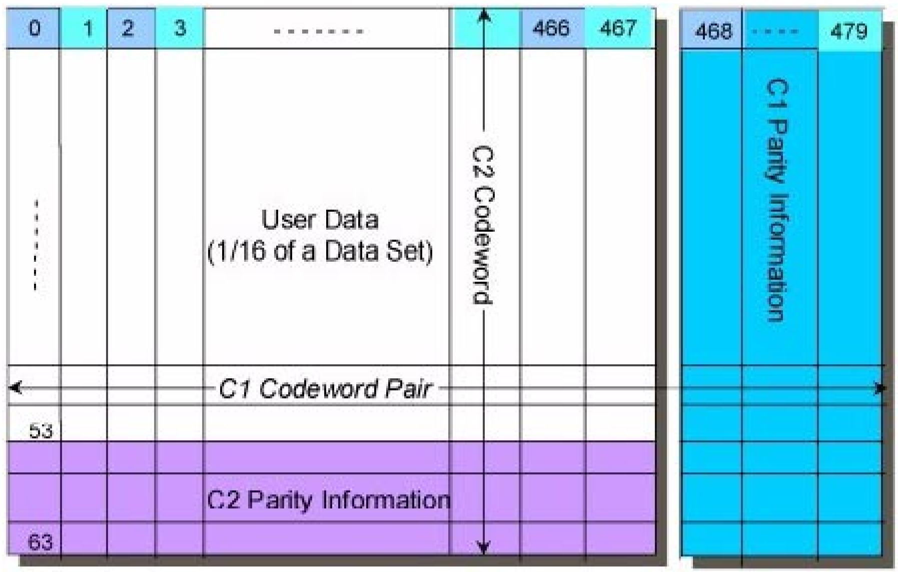

- C1-codeword pair consists of two 240-byte codewords. The first C1-codeword takes each even-numbered byte in the row, and the second C1-codeword takes each odd-numbered byte in the row.
- There are 6 bytes of parity in each C1-codeword. Again the even bytes are in one C1-codeword and the odd bytes in the other.

Formatting for Multiple Channels

The data is then formatted for the multiple head channels:

- Two C1-codeword pairs from two consecutive sub-data sets are placed end-to-end to form a C1-codeword quad (CCQ).
- Eight CCQs are written simultaneously on the eight head channels. This is called a CCQ set and is the smallest unit of data that can be rewritten.
- The timings of the eight channels do not have to be exactly synchronized because each channel has a different and complete CCQ.

The data is then written to tape:

down the tape =>

channels: **0**	0	9	18	27	36	45	54	63	...	475	484	493	502	511
1	1	10	19	28	37	46	55	56	...	476	485	494	503	504
2	2	11	20	29	38	47	48	57	...	477	486	495	496	505
3	3	12	21	30	39	40	49	58	...	478	487	488	497	506
4	4	13	22	31	32	41	50	59	...	479	480	489	498	507
5	5	14	23	24	33	42	51	60	...	472	481	490	499	508
6	6	15	16	25	34	43	52	61	...	473	482	491	500	509
7	7	8	17	26	35	44	53	62	...	474	483	492	501	510

- The number in each box refers to the identity of the CCQ within the data set— there are 512 CCQs in a data set.
- The shaded boxes all contain symbols from a single C2-codeword. C2 codewords are written diagonally across channels.
- A CCQ set consists of the set of CCQs that are written at the same time; this represents a single column in this diagram.

Read-After-Write

After the data is written to tape it is immediately read back to check that it is written correctly.

- Read-after-write detects badly written CCQ sets by checking the C1 ECC code.
- A bad CCQ set is rewritten later down the tape, with the CCQs rotated around the tracks so that no CCQ is written on the same track that it was before.
- A CCQ set is rewritten until at least one good copy of each CCQ is written to tape. (This does not require an entire good CCQ Set).

Definitions

data set: Fixed size block of compressed host information. (404352 bytes)

sub-data set: One sixteenth of a data set (25272 bytes)

ECC-encoded sub-data set: A sub-data set with C1 and C2 protection added

C1-codeword pair: One row of 480 bytes from an ECC-encoded sub-data set

C1 ECC: C1 error correction code is (240,234,7) Reed-Solomon code. Two codewords are interleaved on each row (480 bytes).

C2 ECC: C2 error correction code is (64, 54, 11) code. Each column has one C2-codeword.

CCQ: C1-Codeword Quad. Two C1-codeword pairs, from consecutive sub-data sets

CCQ Set: A number of CCQs that are written simultaneously onto the available channels (8 or 4). This is the smallest rewritable unit of data.

DSIT: Data Set Information Table. Part of the data set that describes its contents.

Servo Format

The Servo Track

The Ultrium format specifies that a servo track will be used to position the head to the required track positions. As a result Ultrium format drives do not have servo write heads but only servo read heads. HP Ultrium drives have 8 servo read heads the layout of which is discussed further in “Head Design”.

Caution Never bulk-erase or degauss Ultrium data cartridges. This will erase the servo information making the cartridges useless.

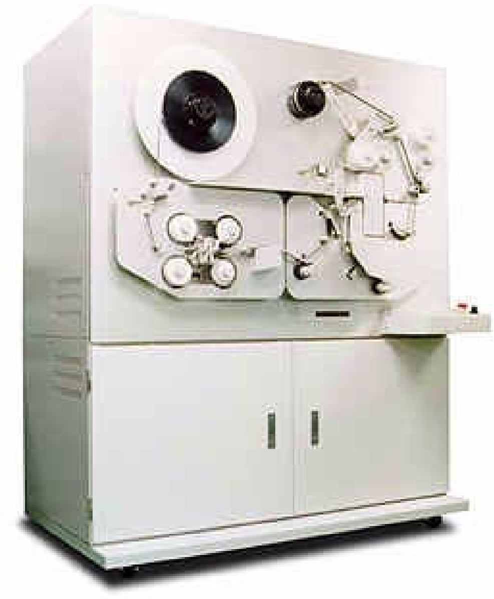

Servo tracks are written during manufacture by an Otari Servo writer, shown on the left. The machine has a very accurate tape path and writes a precise set of five servo tracks down a large spool of tape. The tape is then cut into sections and loaded into Ultrium data cartridges.

The manufacturing information, including the write head assembly identifier, is written into the LTO-Cartridge Memory by the manufacturer and also encoded into the servo track.

When reading or writing data, the servo heads follow the servo bands both above and below the data band. Although only one servo head is required to position the head, the second servo head is there as backup on a different section of tape; it is used if servo is lost with the first head.

The servo band is only wide enough to allow servo following over 6 of the 12 track positions (Gen 1 drives) or 8 of the 16 track positions (Gen 2) in a data band. So that position control is available over all 12 or 16 tracks, there is a second set of servo heads at a different offset relative to the read and write heads. This is shown in more detail in "Head Design" on page 29.

How Servo Following Works

The servo band is much wider than the servo head. The head reads pulses from the servo band and as the head moves up and down over the band the timing of these pulses changes. This allows drive to identify six (Gen 1) or eight (Gen 2) discrete vertical positions. As described above there is enough resolution to provide six or eight positions for writing data.

In addition to allowing accurate vertical positioning, the servo also contains encoded absolute horizontal position data. Each servo frame is 200 microns wide and 36 of these frames make up on LPOS word (taking up 7.2 mm of tape). This word includes an incrementing counter as well as an extra character which over a number of LPOS words builds up to provide manufacturing information about the media.

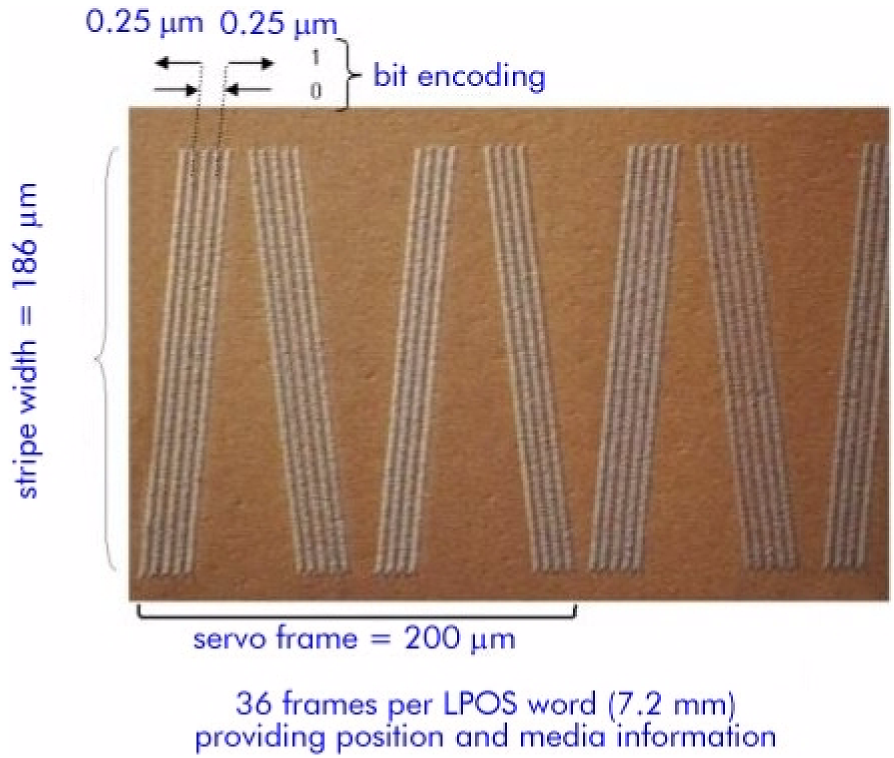

As the photograph above shows, the 2nd and 4th bands of the servo track are shifted relative to the other tracks, in order to be able to write 1 and 0 bits, and hence the position and manufacturer information (LPOS).

Hardware

3

Hardware Block Diagram

Generation 1 Drives

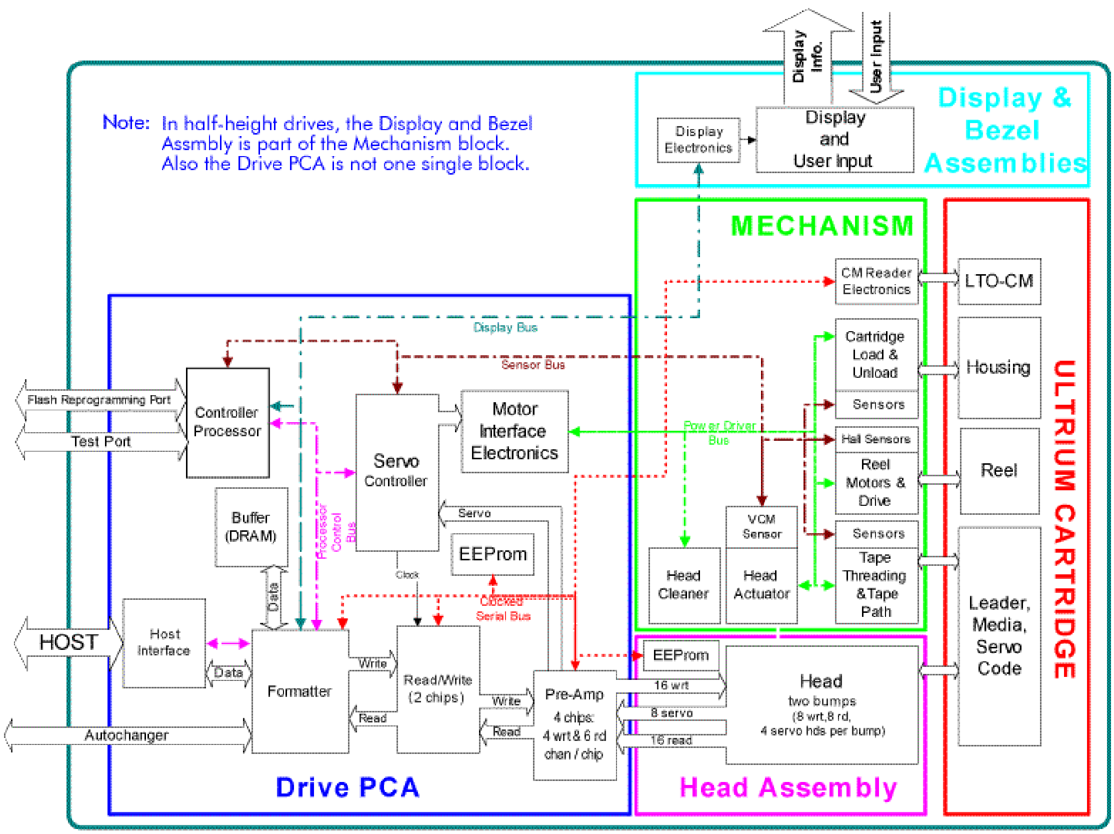

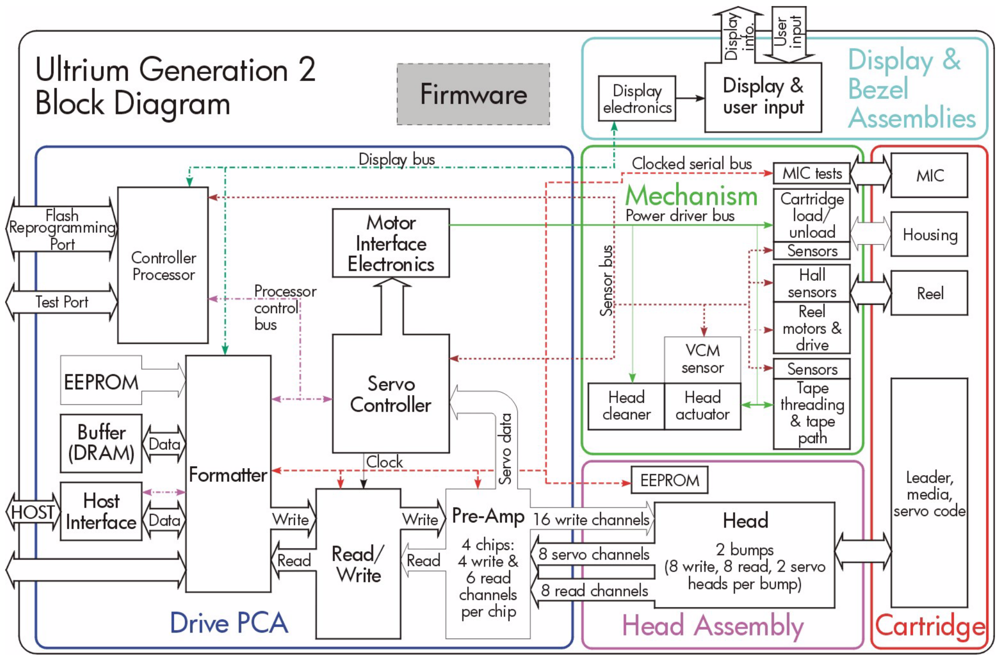

For more details of areas of these diagrams, see the following:

- Cartridge—see page 51
- Controller—see page 46
- Mechanism—see below
- Front Panel—see Chapter 6, "Front Panel LEDs" in **Hardware Integration Guide**, Volume 1 of the HP Ultrium Technical Reference Manual
- Head Assembly/Servo—see page 29

Also see Chapter 2, "Using Special Features," in the **Hardware Integration Guide**, Volume 1 on the HP Ultrium Technical Reference Manual for details of the Automation Control Interface.

Mechanism

Key Features

- Standard 5-1/4 inch full-height form factor
- Soft cartridge load and eject
- A life of 100,000 cycles of load and eject
- Flawless leader threading
- Active servo head positioning
- Air flow management to keep the head and tape cool and clean

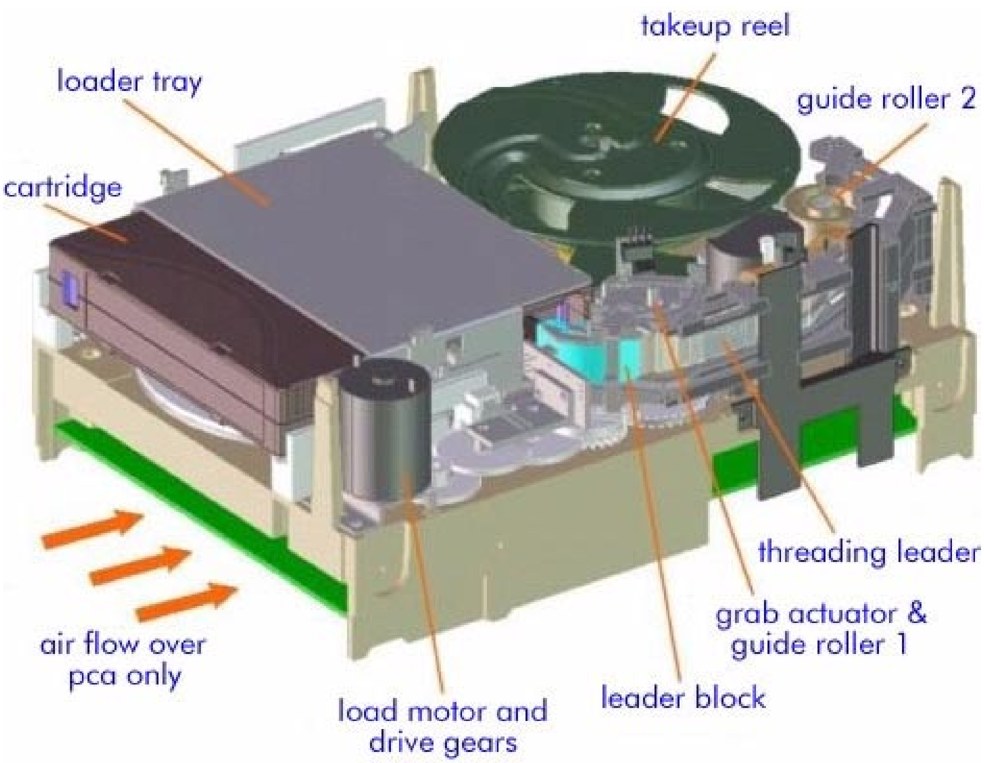

The picture above and the following description refer specifically to a Generation 1 full-height tape drive mechanism, but the principle is the same for half-height and Generation 2 mechanisms.

Hardware

The mechanism is split into two sections by the chassis.

- Above the chassis are the tape handling and tape motion areas, which are designed not to need a cooling airflow.
- The main PCA is mounted underneath the chassis, with vents in the front panel and housing to ensure maximum airflow is drawn over the PCA.

This ensures the maximum cooling for the main PCA and reel motors and the lowest contaminant build up in the tape handling area. Significant testing has been performed on the mechanism to ensure resistance to dust contamination. The drive has an extremely high tolerance to dust contamination.

The mechanism load motor drives the powered load and unloads of the cartridge as well as the leader pin grab mechanics. With a powered cartridge load/unload and leader grab it is important to have a way to extract the cartridge in case the drive power fails. Although not shown in this illustration, one of the gears at the right side of the mechanism can be turned as a thumbwheel allowing manual mechanical load, unload, and leader pin grab/ungrab. This could be useful in the case of a power or PCA failure event with a cartridge loaded. You can turn the supply reel motor manually using a Torx-10 screwdriver through the access points in the mechanism casing. A detailed explanation of this process can be found here.

The HP Ultrium-1 mechanism has no speed encoder as found in DLT tape devices; instead the tape speed is calculated both from hall sensors in the reel motors and also from the servo signal on tape. The servo controller processes these to ensure the tape is kept at the required speed.

Sensors

HP Generation 1 drives have nine and Generation 2 drives have ten sensors. All of them are photo interrupters:

EJ **Eject** position sensor, used to determine when the cartridge has been ejected. It is activated when the cartridge is nearly fully ejected.

CD **Cartridge Down** sensor, used to determine whether the cartridge has been loaded correctly. It is activated when the cartridge is down in its operational position.

CG **Cartridge Grab** sensor, activated when the grabber is fully rotated clockwise, at which point it will have engaged the leader block onto the leader pin assembly.

LP1 **Leader Pin** sensor 1, activated when the top of the leader pin is correctly located in the leader block.

LP2 **Leader Pin** sensor 2, activated when the bottom of the leader pin is correctly located in the leader block.

RD **Rotation Detector** sensor, activated as the load motor moves.

WP **Write Protect** sensor, activated when the cartridge is down and write protect is enabled.

HC1 **Head Cleaner sensor 1**, activated when the head cleaner brush is parked.

HC2 **Head Cleaner sensor 2**, activated when the head cleaner brush is over the heads.

All sensors are enabled during loading, grabbing, threading, ungrabbing or unloading. In order to save power and improve sensor and drive life, no sensors are enabled once the thread has completed.

Soft Loading

- VCR type motorized load mechanism
- Eject distance 25 mm from the front of the form factor (20 mm from the front of the Front Panel)
- 2.5 mm lead-in on all sides

Head Design

Writing Data to Tape in Generation 1 Mode

There are 384 data tracks on a tape, numbered 0–383. They are written in four ‘groups’ called data bands, separated by pre-written servo tracks.

The four data bands are shown between the five servo bands in the diagram below. Each consecutive data band must be filled before the drive starts to write into a different data band. The drive writes the data bands in order 0 through 3. If a tape is damaged, it is normally the edges of the tape that are affected. This is why the first two data bands are in the center of the tape, improving data reliability.

As discussed in "The Servo Track" on page 22, the servo heads and servo tracks allow the head to be positioned at 12 unique positions within each data band. With an 8-channel head this allows 96 tracks to be written within a data band. As there are four data bands, a total of 384 tracks can be written onto the tape.

- Data is written to the center of the tape first, then nearest the Reference Edge, so data bands are written in the order 0–3.
- Each data channel writes 12 tracks of data in each data band. They are written in the order shown to the left starting at 0.
- The tape is filled up by writing 48 'wraps' (transferring tape from one reel to the other and then back again).

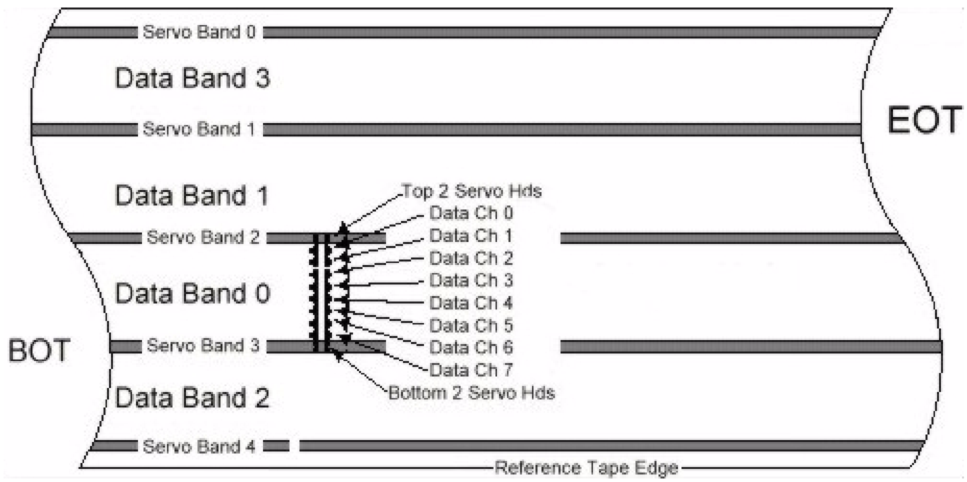

Within the data band the head serpentines, completing a full travel to the end of the tape before changing the head position and writing a track in the reverse direction back down the tape.

The diagram below shows the order in which the drive writes data tracks within the data band.

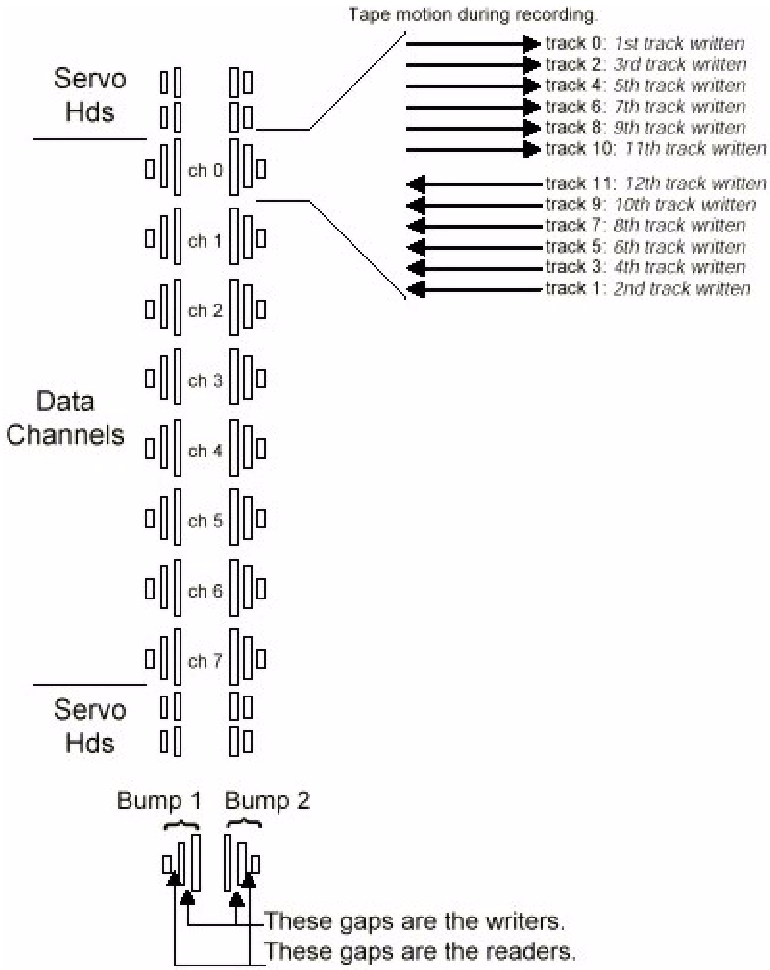

To allow read-after-write, the data is always written with the 'leading' bump:

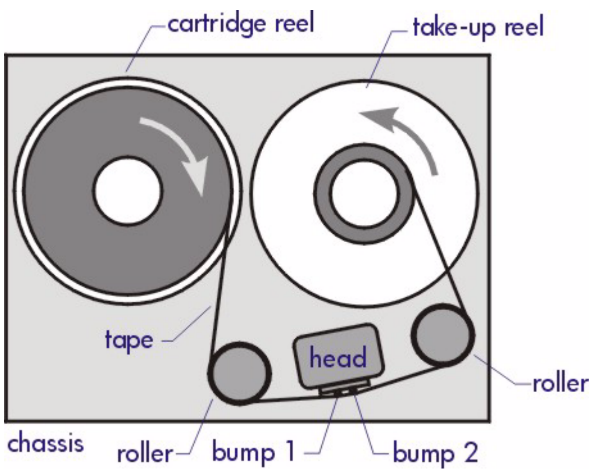

The rotation shown is when reading or writing even-numbered tracks. For odd-numbered tracks, both reels rotate in the opposite direction.

Even-numbered tracks:
Head bump 1 writes
Bump 2 performs read-after-write.

Odd-numbered tracks:
Bump 2 writes.
Bump 1 performs read-after-write.

Writing Data to Tape in Generation 2 Mode

There are 512 data tracks on a tape, numbered 0–511. As in generation 1 drives, they are written in four 'groups' called data bands, separated by pre-written servo tracks. In generation 2 drives, the servo heads and servo tracks allow the head to be positioned at 16 unique positions within each data band. With an 8-channel head this allows 128 tracks to be written within a data band. As there are four data bands, a total of 512 tracks can be written onto the tape.

- Data is written to the center of the tape first, then nearest the Reference Edge, so data bands are written in the order 0–3.
- Each data channel writes 16 tracks of data in each data band. They are written in the order shown to the left starting at 0.
- The tape is filled up by writing 64 'wraps' (transferring tape from one reel to the other and then back again).

Within the data band the head serpentines, completing a full travel to the end of the tape before changing the head position and writing a track in the reverse direction back down the tape.

The diagram below shows the order in which the drive writes data tracks within the data band.

To allow read-after-write, the data is always written with the 'leading' bump:

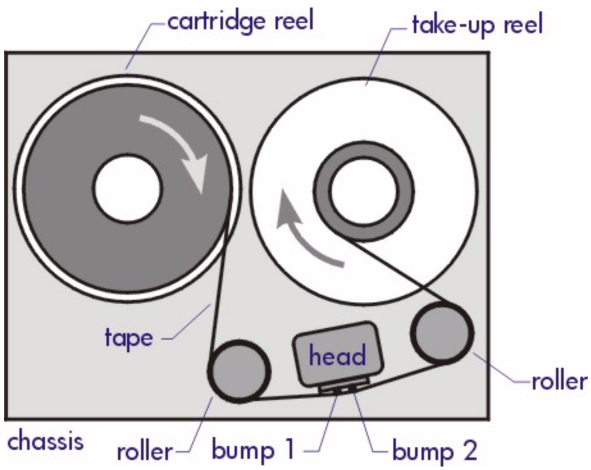

The rotation shown is when reading or writing even-numbered tracks. For odd-numbered tracks, both reels rotate in the opposite direction.

Even-numbered tracks:
Head bump 1 writes.
Bump 2 performs read-after-write.

Odd-numbered tracks:
Bump 2 writes.
Bump 1 performs read-after-write.

Head Design

In order to be able to write at high transfer rates, an 8-channel read write head is used in HP Ultrium drives. The head assembly is 'printed' onto a ceramic substrate using photolithography techniques. HP uses head assemblies from STK and Seagate.

The heads in HP Ultrium drives are magneto resistive (MR). MR refers to the semiconductor used to read magnetic information from tape, which changes its resistance according to the magnetic field put across it. The head writes magnetic information onto tape using a copper coil that is 'grown' onto the ceramic element (sometimes referred to as thin film).

The head assembly is an 8-channel head, but consists of 16 read and 16 write elements shared between two 'bumps' that allow data to be written in both forward and reverse directions down the tape.

As can be seen in the diagram below, each bump has:

- 2 servo elements at the top of the bump
- 2 servo elements at the bottom of the bump
- 8 read elements

- 8 write elements

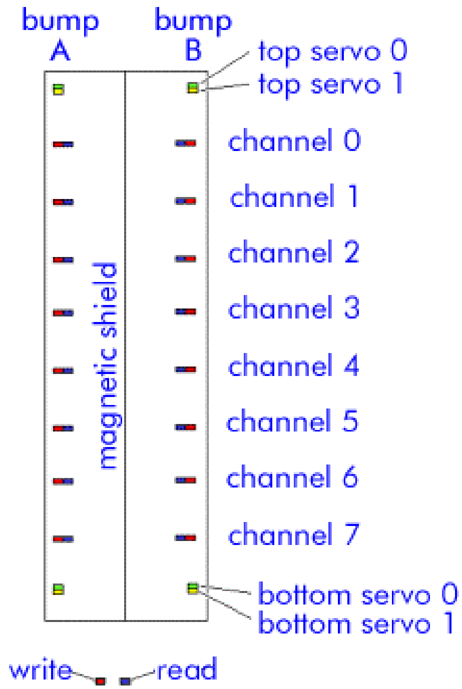

The illustration below shows the actual head assembly. The leading edges are designed to wipe any debris from the tape before it gets to the heads. The tape bearing surfaces are ceramic to resist wear so that the head has a life of >30,000 hrs.

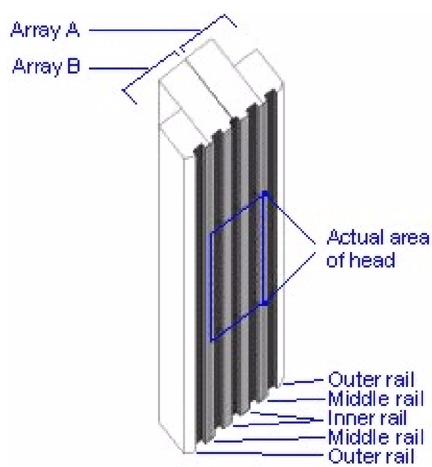

Data is always written with the leading bump:

- If the tape is traveling from left to right (forward) data is written with Bump A and read-after-write with Bump B.
- If the tape is traveling from right to left (reverse) data is written with Bump B and read-after-write with Bump A.

It is not practical to perform read-after-write with the read head on the same bump as the write head. This would produce significant crosstalk between the signals from the heads, which would result in a high error rate.

Head Cleaning

There are two mechanisms used for cleaning the drive in the HP Ultrium tape drive. One is the internal head cleaning brush within the drive, the other by inserting a cleaning cartridge and pulling cleaning tape through the tape path.

Internal Head Cleaning Brush

The internal head cleaning brush will be activated if one or more of the following is met:

- 200 km of tape has been pulled across the heads (equivalent to using five full Generation 2 data cartridges).
- This is the 25th cartridges that has been unthreaded since the last cleaning.

Once the head cleaner has been triggered by either the tape counter or the cartridge counter, both are reset. It will then take another 200 km tape pulling or another 25 cartridges unthreaded before the next head scrub becomes due. Both of these counters are also reset at power-on.

Cleaning Cartridge

The tape drive tells the automation controller that a cleaning tape needs to be used through two bits in the ACI `get_drive_status` command. The Cleaning Needed bit indicates deterioration in the write or read margin of the drive and hence it is recommended that a cleaning cartridge be inserted into the drive at the earliest opportunity. Following a successful clean, the Cleaning Needed bit will be cleared. The Cleaning Required bit indicates that the drive is unable to read or write unless the drive is first cleaned with a cleaning tape. It is recommended that a cleaning cartridge be loaded into the drive immediately. Following a successful clean, the Cleaning Required bit will be cleared.

Note that cleaning cartridge use has been extended for Gen 2 Ultrium drives from 15 to 50 cleanings.

Handling Cleaning Cartridges

The actions that the drive takes when a cleaning cartridge is inserted are as follows:

1 The cleaning cartridge is drawn into the drive so the Cartridge Memory can be read. The Cartridge Type field in the Cartridge Memory is used to detect that it is a cleaning cartridge.

2 The cartridge is considered to be compatible with the drive if one of the following applies:

   — The Cartridge Memory contains Drive Manufacturer Support pages and 'HP' is listed in them.

   — The Cartridge Memory has no Drive Manufacturer Support pages but the Cartridge Manufacturer field in the memory contains a label relating to an approved cartridge manufacturer.

   If the cartridge is considered incompatible with the drive then:

   a The cartridge is ejected.

   b The Tape Error LED is lit indicating that the tape is defective. A CHECK CONDITION occurs with sense key of Medium Error and additional sense of 3007h (cleaning failure).

   c The Invalid TapeAlert Flag (23) is set.

   d If the Cleaning LED and TapeAlert flag 20 were set before the cleaning cartridge was inserted, they remain set.

   e Cartridge handling is terminated.

6 Checks are made to avoid excessive cleaning cartridge usage. To prevent damage to the heads, the cleaning tape is only used if at least one of the following is true:

   — The drive is indicating that it needs to be cleaned (the Cleaning LED is on and TapeAlert flag 20 is set).

      — A cleaning cartridge has not been used in the last 336 hours (2 weeks) since the drive was powered on.

Instead, the procedure is as follows:

a The internal head cleaner is operated.

b The usage tracking recorded in the Cartridge Memory is not incremented.

c Cartridge handling is terminated.

Note On full-height drives, if the drive is in thermal failure mode when a cleaning cartridge is inserted, the internal head cleaner is operated and the cartridge ejected without using the cleaning tape.

4 The cartridge is checked to see if has expired.

The thread count and tape length in the Cartridge Memory are checked to see if the cartridge has been used the maximum number of times. If it has, the procedure is as follows:

a The cartridge is ejected.

b The Tape Error LED is lit indicating that the tape is defective. A CHECK CONDITION occurs with sense key of Medium Error and additional sense of 3007h (cleaning failure).

c The Expired Cleaning Media TapeAlert Flag (22) is set. If the Cleaning LED and TapeAlert flag 20 were set before the cleaning cartridge was inserted, they remain set.

d Cartridge handling is terminated.

5 The cleaning tape is used to clean the heads. (See below for the detailed procedure.)

6 If necessary, the Cleaning LED is switched off and TapeAlert flag 20 is cleared.

7 The cartridge is ejected.

Cleaning Procedure

1 Thread the tape.

2 Move to the starting point on tape while wiggling the heads.

Each clean uses 18.4m of tape. The first clean is started at 8m so the starting point is [8 + 18.4 * (thread count - 1)]m from BOT.

3 Run the tape over the heads at 4 m/s for 18.4m while the heads wiggle.

4. Run tape over the heads at 4 m/s for 18.4m in the reverse direction while the heads wiggle.
5. Rewind the tape.
6. Unthread the tape.
7. Operate the internal head cleaner.

Using the Internal Cleaner

The internal head cleaner is a brush that can be drawn over the heads to 'scrub' them. It is effective at removing loose debris. As it is less abrasive than cleaning tape it can safely be used more frequently.

The internal head cleaner is used in the following circumstances:

- Immediately after using a cleaning tape.
- Instead of the cleaning tape if a cleaning cartridge is used too often.
- As part of error recovery.
  If multiple read retries have failed, the tape is unthreaded, the head cleaner is operated, the tape is re-threaded, and retries begin again.
- After an unthread if more than 150,000 metres of tape have been pulled since the last use of the head cleaner. (This is equivalent to slightly more than five full uses of 200 GB Ultrium 1 cartridges)

Leader Threading

In Ultrium drives, reliable threading is achieved through the following measures:

- A hybrid solution using a threading leader and rigid parts
- A positive snapping feature
- A mechanical interlock to prevent the leader being swallowed
- Redundant sensors to detect the proper connection and prevent leader loss
- Space efficiency for a small form factor

With a single reel tape device, the interlock between the take-up reel and the media is the most important part of the drive for reliability.

As specified in the Ultrium format, with HP Ultrium drives, there is a leader pin attached to the end of the tape.

You can easily check the leader pin in a cartridge by sliding open the cartridge door and checking that the pin is seated as shown in the diagram below.

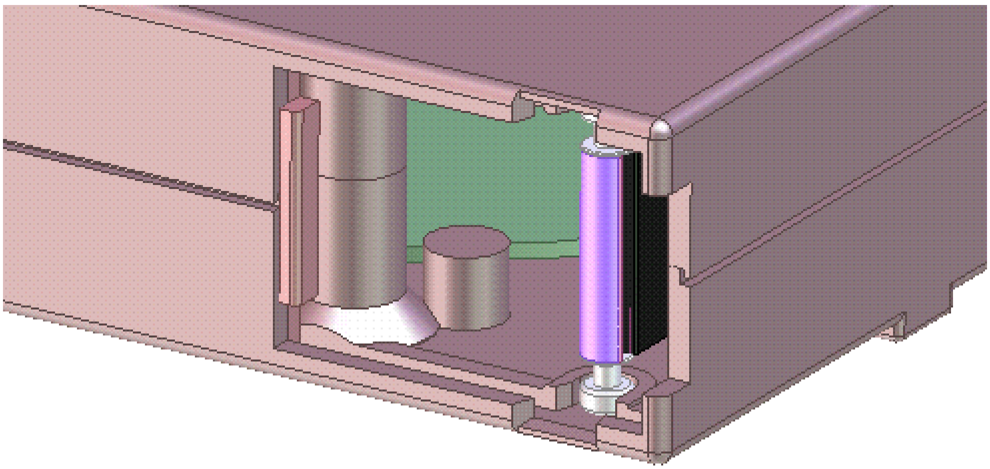

Within the drive, a leader block engages directly onto the leader pin at the end of the tape. It has a spring retainer to maintain the grab in case of power failure or loss of tension.

Opto-mechanical sensors at the top and bottom edges of the leader block confirm that the leader block is engaged properly and then the block is pulled through the drive and onto the take-up reel.

The leader block is designed so that it completes the spindle of the take-up reel.

Previous tape technologies have fixed the position of the tape during motion (called the tape path) very precisely and used a stepper motor to move the head across the tape. HP Ultrium drives use an active servo system to move the head rapidly in response to any vertical movement of the tape. Because of this, the two guide rollers do not need to 'pinch' the tape into an exact position but allow the tape to float causing less wear to the media. See "The Servo Track" on page 22.

With any linear tape device there is a potential for problems caused by rollers damaging the tape and from worn rollers distorting the tape path. Because the tape is not pinched by the rollers, edge wear is significantly reduced. In addition the rollers have a wear-resistant coating to prevent any channels being eroded into the surface. HP uses an SEM to check that roller wear is within acceptable limits.

Controller

4

The controller contains all the hardware and firmware that communicates via the external interfaces and controls the device electronics.

Architecture

The controller is based on high performance R/W pipeline,which includes in Generation 1 drives a 16 MB buffer and in Generation 2 drives a 64 MB buffer.

The SCSI and FC host interfaces are interchangeable.

- All external interfaces share some common commands
  - Host, Front Panel, Automation, Test Port
- Partitioned into functional modules with assigned owners.
- Designed with future generations in mind.

For an overview of the logical format see Chapter 2, “Format” in **Background to Ultrium Drives**, Volume 6 of the Ultrium Technical Reference Manual

Hardware

The following picture is of the Controller PCA for Generation 1 full-height drives only:

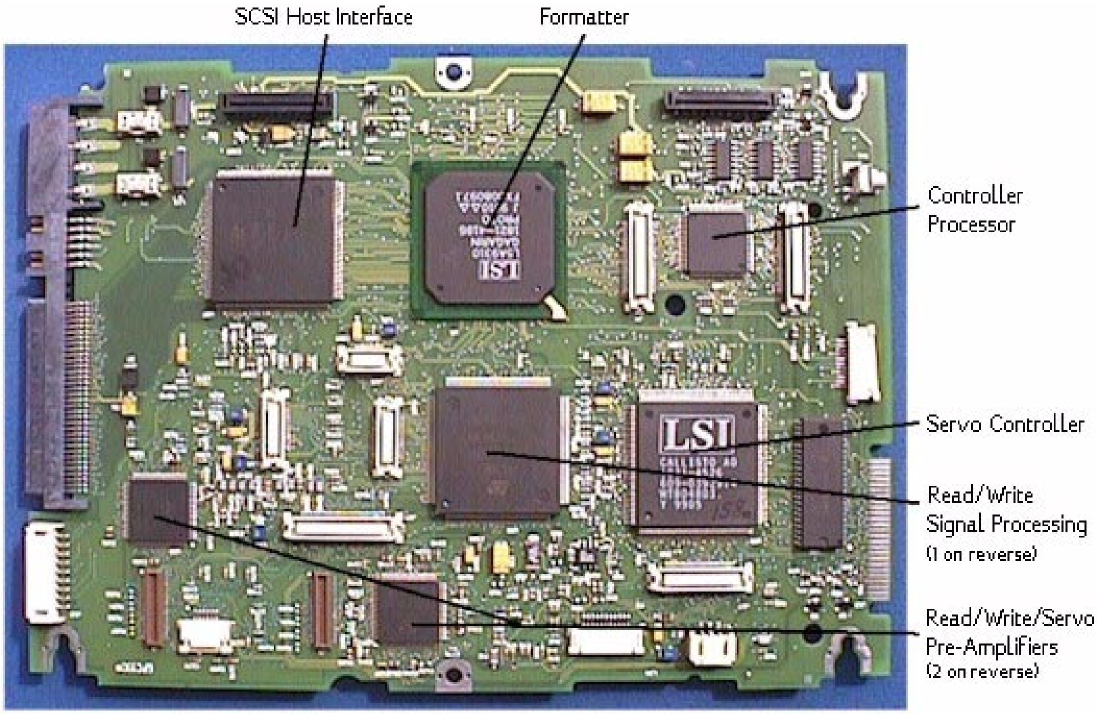

The following picture is of the Controller PCA for Generation 2 full-height drives only:

SCSI Host Interface ASIC

There are two ASICs that may be mounted, one for SCSI, the other for FC.

Generation 1 drives: The ASIC also generates or checks CRC (Cyclic Redundancy Codes) on each record as required by the Ultrium tape format specification.

Generation 2 drives: The SCSi ASIC has an integrated ARM9 microprocessor. The FC ASIC has an integrated ARM9 microprocessor. The firmware that runs on the ARM controls SCSI/FC activities above the SCSI/FC phase level. As a result, this is highly flexible and maintainable.

Formatter ASIC

The Formatter ASIC contains data compression, ECC, channel mapping and RLL encoding/decoding functions to meet the Ultrium tape format specification for the 8-channel format. In addition, it contains the interface logic for the serial devices on the product (EEPROMS, LTO-CM device reader, Read/Write Signal Processing ASICs and Read/Write/Servo Pre-Amplifier ASICs), the front panel graphics interface and the autochanger/library interface. Finally, the device contains the Buffer Manager for the 16/8 MB data buffer (Gen 1 drives) or 64 MB buffer (Gen 2 drives). Gen 1 drives: This device is a 272-pin BGA device that operates at 66 MHz (derived from an 13.2 MHz clock) and with a 3.3 volt supply. *Gen 2 drives:* This device is a 456PBGA package, operating at 132 MHz; it has a 1.8V core and 3.3 volt I/O.

Controller Processor ASIC

The Controller ASIC contains the microprocessor (NEC V851 core), RAM and flash ROM on which the product firmware resides and executes. The firmware manages and controls the operation of the drive.

The ASIC is a 100-pin TQFP device that operates at 33 MHz (based on a 6.6 MHz clock) and with a 3.3 volt supply. It can be flash programed; control is over a 7.8V Vpp line and the code is sent through a clocked serial port.

Servo Controller ASIC

The Servo Controller ASIC provides the closed loop control of the vertical position of the head/actuator mechanics. Within the ASIC is a DSP core, analog to digital and digital analog converters, servo signal demodulator functions and logic to interface to the mechanism interface electronics. In addition, the ASIC provides the master read and write clock for the system so that data that is written or read tracks the tape speed. This device is a 208-pin PQFP device that uses a 50 MHz clock and with a 3.3 volt supply.

Read/Write Signal Processing ASIC *(Gen 1 drives)*

The Read/Write Signal Processing ASICs provides the functions to filter and peak-detect the read data from the pre-amplifiers. They provide the resulting synchronized data and clock to the Formatter ASIC. The ASIC also takes the write data and clock, and generates the write-equalized transitions as stipulated in the Ultrium tape format specification. Each Read/Write Signal Processing ASIC has four parallel channels; two are required. This device is a 144-pin PQFP device that operates with a 5 volt supply.

Read Analog Front-End (AFE) ASIC *(Gen 2 drives)*

The Read Analog Front-End (AFE) ASIC is an 8 channel LTO-1/2 mixed-signal ASIC, taking 8 analog channels from the head pre-amp. Each analog channel is fed to an amplifier with Automatic Gain Control. A Continuous Time Filter provides for equalization of the signal. Finally, the signal is fed to an ADC for quantizing into an 8-bit quantity. The eight 8-bit buses are multiplexed and fed to the Digital Back-End (DBE) for further signal processing. Quantization is performed at a slightly higher rate (1.25x to 1.5x) than the actual encoded bit rate within the data stream. This provides an SNR advantage and simplifies the analog processing requirements.

This device is a 176-pin BGA fabricated in 0.18μm 5-layer metal CMOS process, requiring 3V3 for the IO and 1V8 for the internal design.

Read Digital Back-End (DBE) and Write Path ASIC *(Gen 2 drives)*

The Read Digital Back-End (DBE) and Write Path ASIC is an 8 channel LTO-1/2 fully digital signal processing ASIC, taking 8 quantized channels of 8-bit data from the AFE and processes them to extract recovered data. Each channel is fed to an asymmetry correction unit (to eliminate asymmetry introduced by the non-linear features of the head), then onto a filter to provide extra equalization. From there it enters a digital PLL to extra a bit clock and interpolated data. This interpolated data is fed to an adaptive filter to increase SNR and finally the data is fed to Viterbi-like decoder to extract individual bits. 8 channels of recovered data are fed to the formatter ASIC which can then extract user data. The DBE contains a huge amount of metric collection devices to allow the signal through the ASIC to be monitored and characterized. Features include a histogram, a correlator and a SOVA block.

This ASIC also contains a WPC/WEQ block so that write data from the Formatter ASIC can be equalized and pre-compensated before being applied to the write amps within the Pre-Amp ASIC. This device is a 320-pin BGA fabricated in 0.18um 5layer metal CMOS process requiring 3V3 for the IO and 1V8 for the internal logic.

Read/Write/Servo Pre-Amplifier ASICs

The Read/Write/Servo Pre-Amplifier ASICs provide the read pre-amplification function for the data and servo heads. In addition, they provide the write drive function for the data heads. The ASICs are 100-pin TQFP devices that operate with both the 3.3 volt and 5 volt supplies. Each Read/Write Signal Processing ASIC supports four data channels (write and read) and two servo channels (read only). Consequently, four Read/Write/Servo Pre-Amplifier ASICs are required to interface with the head assembly.

Motor Interface Electronics

The motor interface electronics provide the appropriate signal conditioning and power amplification to drive the motors of the mechanism. In addition, they provide the drive for the actuator VCM, and the interface conditioning and switching logic to the opto sensors on the mechanism. Given the varied nature of these functions, they are implemented on the PCA using discrete components and “off-the-shelf” integrated circuits.

Buffer

The function of the buffer is to provide temporary storage capacity so that drives can hold user data during repositions and turn around without having to pause the data transfer from the host. The buffer size in Generation 1 drives is 16 MB organized as 1M x 16 bits x 4 banks SDRAM. In Generation 2 drives it is 64 MB organized as 4M x 16 bits x 4 banks SDRAM.

EEPROM

The EEPROM on the PCA is used to store important product parameters and logs that are required to be maintained through a power failure. The EEPROM is implemented in Generation 1 drives using a 64 KB SPI serial EEPROM. In Generation 2 drives it uses a 256 KB SPI serial EEPROM.

Firmware

The firmware is developed in high level ‘C’ and SDL languages, and implemented using a ‘state of the art’ CASE tool. It executes on the NEC V851-based Controller Processor ASIC, which has the following features:

- 32/16 V850 core
- 336 KB of flash program memory

- 26 KB of SRAM
- Operation at 33 MHz

Firmware can be downloaded into the drive in various ways:

- via the Host Interface (SCSI or FC)
- via a tape cartridge

Modular View

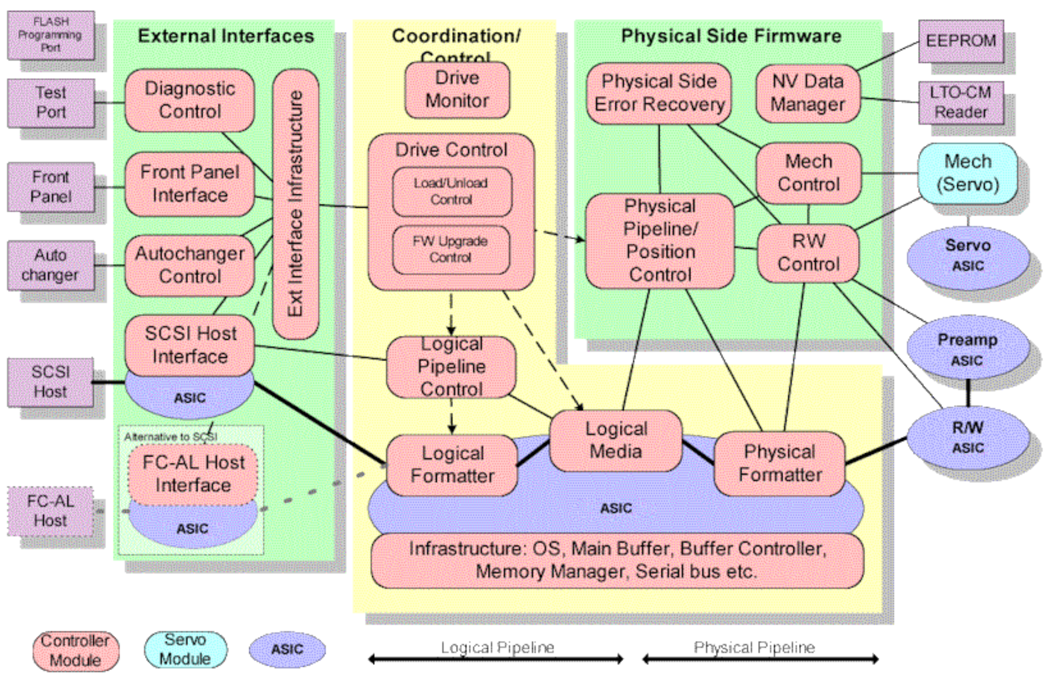

The diagram can be divided into a number of areas:

- At the bottom is the infrastructure, providing the operating environment for the modules
- Above the infrastructure is the main data path (shown with a thick black line) made up of a logical pipeline (host focussed) and a physical pipe-line (tape focussed).

- Above the data path are the logical and physical pipeline controllers
- Above the pipeline controllers is Drive Control, in charge of the whole drive.
- Between the two pipelines is the Logical Media which is the main data buffer for the drive.
- On the left are the four external interface modules taking command input from the four external user types. These use a local 'External Interfaces infrastructure' for passing command between modules.
- On the right are the modules Mech and Read/Write with which the controller communicates, but which are not considered to be part of the controller.

Briefly, the modules have the following functions:

- **Drive Control** — the dominant module. Drive Control has ultimate control over the whole drive from a module point of view. It can pass control to Diagnostics Control but can get it back when it needs to. It also manages the load and unload processes and the reprogramming of the firmware.
- **Diagnostic Control** — a major module. Diagnostic Control has a port that is continuously live, allowing debug, manufacturing (MTC — Manufacturing Test Code) or field access to monitor, control or test the internals of the drive. It executes most commands locally or it can pass them on to other modules (for example, a mechanism pass-through command). The module can operate either at the module level or the block level (that is, the ASICs, DRAMs and so on). It needs permission from Drive Control to take control before it is allowed to change the behavior of any of the blocks or modules.
- **Logical Pipeline Control** — a major module. This controls the use of the logical pipeline for all data transfers between the Host Interface and the Logical Media.
- **Physical Side Error Recovery** — a minor module. This controls the whole physical pipeline if error recovery is required that is beyond the capability of the physical transfer and position controllers. The architecture offers maximum error recovery capability without specifying the algorithm.
- **Physical Pipeline Control** — a minor module. It has overall control over the physical pipeline during data transfers and positioning operations coordinating the tape position and speed with the actual formatting of the data. It is also responsible for determining the optimum tape speed for host throughput matching and minimum head/tape/mechanism wear.

- **Mech Control** — controls the mechanism (tape motion motors, head actuator movement, cartridge movement, physical sensors and so on) via the Mech interface (shown as a non-controller module in the diagram). Note that although the Mech interface is external to the controller architecture it does have a module interface and therefore makes use of the infrastructure (and the Controller CPU).
- **Host Interface** — manages the host interface protocol (SCSI or FC) and executes host commands locally if it can. It passes non-data path commands (such as load tape) to Drive Control, requests for transfer modes (such as streaming writes) to the Logical Pipeline Controller and data (records and filemarks) directly to the Logical Formatter.
- **Logical Formatter** — performs logical formatting (data compression, embedding record boundaries and filemarks, creation of Data Sets) on data from the Host Interface and passes completed Data Sets into the Logical Media. It also unformats the data on the way back.
- **Logical Media** — The logical representation of the data on the tape, held in terms of a cache of Data Sets. To the logical pipeline it is the tape. To the physical pipe-line it is the host. It provides host data rate matching by managing the amount of space/data in the main buffer. It allows the host to write, read and space within the buffer without the need for tape movement. Note that in some representations of the architecture, the main buffer is considered part of the infrastructure rather than being internal to the Logical Media.
- **Physical Formatter** — performs physical formatting (generating ECC redundancy, dividing into CCQs, adding headers and so on) on Data Sets from the Logical Media to the Read/Write interface. It also performs error detection and correction on the data on the way back.
- **Automation Interface** — interfaces an Autochanger/Library Controller (not part of the architecture). The Autochanger Controller can be subservient to the drive (as in HP DDS autochangers) or can be in control if the Host Interface is actually part of the Autochanger sub-system (as in HP Storage Library products).
- **Read/Write Control** — controls the set-up of the heads to give optimum read/write performance. It also calibrates the heads during reading and writing by coordinating tape motion through Mech Control and the set-up of the head controllers through Read/Write.
- **Drive Monitor** — monitors the logs held in various modules around the architecture (mainly on the physical side). As a result, it has the best

knowledge of the general health of the drive. It warns Drive Control of any health scares. Tape Alert is implemented in this module.

- **NV Data Manager** — provides access to the non-volatile data within the architecture:
  - The LTO-CM (Memory In Cassette) keeps details of the media such as type and usage.
  - The EEPROM keeps details of the drive, such as manufacturing information, SCSI address, calibration information and so on. There may be more than one EEPROM.
- **Infrastructure** — provides hardware and firmware services for the architecture, such as the operating system, memory manager, trace logs, serial bus, access to main buffer memory and so on.

Cartridges

5

The Ultrium cartridge is a single-reel cartridge containing half-inch tape. At 2:1 data compression, the capacity for a Generation 1 cartridge is 200 GB and 400 GB for a Generation 2 cartridge. Generation 2 drives can read and write both Generation 1 and Generation 2 tapes.

Of particular note are the following features:

- LTO-CM (LTO Cartridge Memory). This is an EEPROM contained within the cartridge shell which can be read by the drive without electrical contact.
- Secure parking of the tape leader pin makes for a robust cartridge.
- Standard automation notches allow cartridges can be used in automated devices such as autochangers.
- Compact, efficient reel locks prevent tape movement when the cartridge is not in a drive.

The size and shape of Ultrium media is very similar to that of DLT cartridges in order to make it easy for automation/library vendors to integrate Ultrium into existing DLT libraries. Many of the automation notches are identical to the DLT IV cartridge so that the same picker designs can be used with both kinds of cartridges.

The photograph below shows a 'certified' Ultrium cartridge carrying the Ultrium logo. All drive and media manufacturers have to submit their products to the LTO CVE (Compliance Verification Entity), which ensures that all drives and media are written to the correct LTO standard. The CVE is at Measurement Analysis Corporation, Torrance, California. Manufacturers can only use the LTO logo once their products have passed this testing.

200 GB* tape length	580m
Tape thickness:	8.9 microns
Tape width:	12.65 mm
Data tracks:	384
Recording Density:	4.88 Kb/mm
Tape Speed:	4.1 m/s
Media life:	1,000,000 passes

*capacities at 2:1 data compression

Cartridge Reliability

Reliability of Ultrium cartridges is enhanced through the following features:

- A compact reel lock, based on 3590-style reel locks
- Positive retention of the leader pin assembly
- Strain-relieved attachment of the tape to the leader pin
- Simple and reliable door mechanism
- No abrasive debris from glass-filled case materials
- A package capable of withstanding reasonable use and misuse
- Specified to 20,000 load and unloads to make cartridges suitable for use in automated devices

Cartridge Features—Top

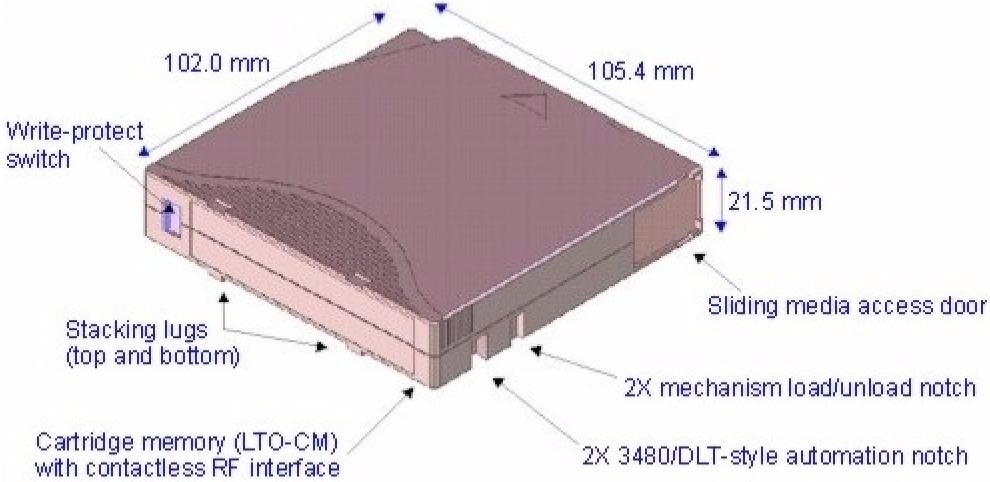

- The stacking lugs allow cartridges to be stacked securely on top of each other.
- The sliding media access door protects the tape from dust and debris when not in a drive.
- The mechanism load/unload notches are used by Ultrium drives to ensure positive and accurate loading and unloading.
- The automation notches are for use in automated devices that use robotics to load catridges into drives automatically.
- For details of the LTO-Cartridge Memory, see "LTO Cartridge Memory" on page 55.

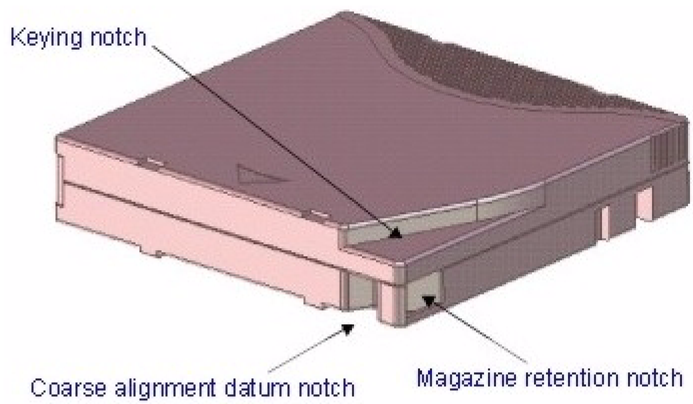

These notches ensure the proper alignment and retention of the cartridge within the drive.

Cartridge Features—Bottom

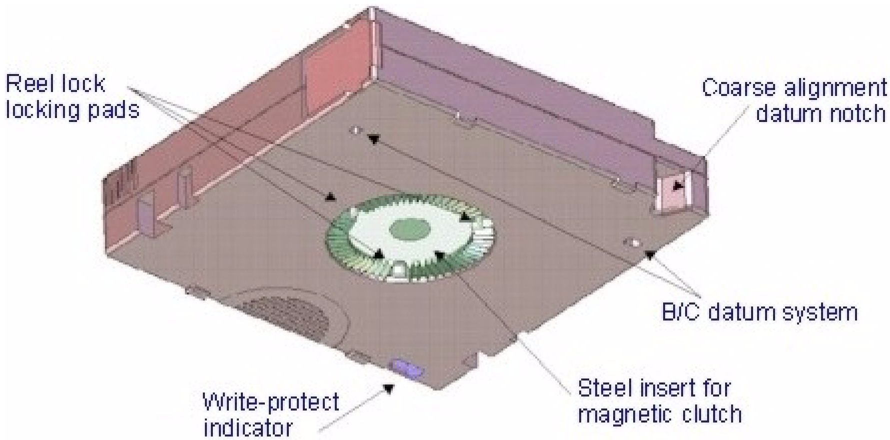

The reel lock locking pads ensure that the reel does not move when the cartridge is not in a drive.

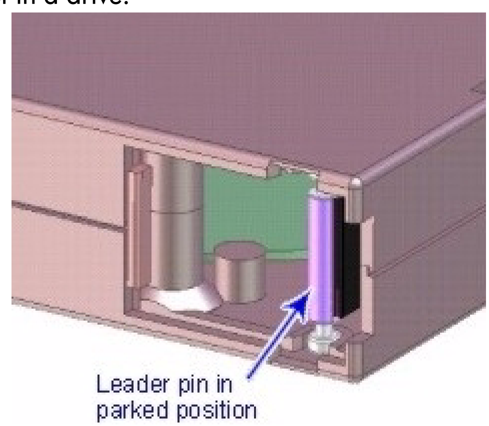

See the “Leader Threading” on page 38 for details of capturing this leader pin.

LTO Cartridge Memory

Linear Tape Open—Cartridge Memory (LTO-CM) is EEPROM memory that is embedded in every LTO Ultrium tape cartridge. It is non-volatile and is contactless in that it is read by inductive coupling rather than electrical contact.

Cartridge Memory has been added to the LTO cartridge for the following reasons:

- It speeds up load/unload times—there is no need to read system areas.
- It speeds up movement around the tape by storing the tape directory (physical to logical mapping).
- It increases tape reliability because fewer tape passes are needed.
- It stores diagnostic and log information for tracking purposes.

Most of these uses will be invisible to applications and handled internally by the drive. There is potential for some other applications which is being investigated; these would use the “Application Specific Data” area within the cartridge memory.

The memory is primarily designed to speed up internal operations in the drive such as loading and spacing, but it also contains free space that can be used by application software. Of the 4 kilobyte memory, about 1 kilobyte is free space. This may be used to store “common” information (shared by all software vendors) and “vendor-unique” information (specific to the application).

A SCSI access method has been defined to allow hosts to use this free space; new Write Attribute and Read Attribute commands have been proposed to ANSI in document T10/99-148. In the proposal, “common” attributes include items like “Application Name” and “Date Last Written”.

There are three types of attributes:

- Medium attributes are hard coded in the Cartridge Memory during manufacture and store such data as a unique serial number.
- Device attributes are maintained by the host and store such data as load counts.
- Host attributes are maintained by the drive and contain such information as a tape label or backup date.

Specification of the EEPROM

- 32 kilobits (4 kilobytes), organized as 128 x 32-byte blocks
- An operational life of >500K write cycles with a 20-year data retention life
- Contains both read-only and rewritable areas
- Word-wide (2 bytes) or block-wide (32 bytes) write size
- Data-protection on write and read with parity and CRC

Interface Specification

- Contactless, passive RF interface using a proximity inductive coupling with a range in the order of millimetres.
- Power to the transponder is coupled through the interface
- The range depends on implementation (10 to 20 mm is maximum). The best error rate performance will occur at short distances
- The memory can be read from below (by a drive) or from the front (in libraries).

Use in Libraries

For suggestions of how to make use of cartridge memory in libraries, see “Using Cartridge Memory (LTO-CM)” in Chapter 2 “Using Special Features in Libraries” of the **Hardware Integration Guide**, Volume 1 of the HP Ultrium Technical Reference Manual.

Features of LTO-CM

- 32 kilobits (4 kilobytes) EEPROM
  - Organized as 128 x 32 byte blocks.
  - It has a life of >500K write cycles and a data retention life of 20 years.
  - Defined with both read-only and rewritable areas.
  - Write size is word-wide (2 bytes) or block-wide (32 bytes).
  - Data is protected on write and read with Parity and CRC.
- Contact-less, passive RF interface
  - A proximity inductive coupling has a range of a few millimeters.

- Power to the transponder is coupled via this interface.
- The range depends on the implementation, with a maximum of 10 to 20 mm.
- The best error rate performance will occur at short distances.

- Based on a proposed standard (ISO 14443)
- Standard across every Ultrium and Accelis cartridge
  - Multiple reading angles in Ultrium cartridge

Contents of LTO-CM

- Read-Only Section
  - CM Transponder manufacturer information (32 bytes)
  - Cartridge manufacturer information (64 bytes)
  - Media manufacturer information (64 bytes)
- Re-writeable Section
  - Initialization data, written when the format is initialized or re-initialized (64 bytes)
  - Usage information (256 bytes)
  - Tape directory (1536 bytes)
  - OED information (64 bytes)
  - Mechanism manufacturer information (~448 bytes)
  - Application Specific Data (~1504 bytes).

Each sub-section is wrapped with CRC for added data integrity.

The Tape Label, Media Manufacturer Information and Initialization Data are also copied to the tape to enable data still to be read if the cartridge memory fails.

More Information

- The latest version of the specification is here. This is nearing agreement.
- The access specification can be found at `http://www.t10.org/doc99.htm`. At the website, click on the version of document 99-148 with the highest "R"number.

Cartridges

glossary

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.

ANSI American National Standards Institute, which sets standards for, amongst other things, SCSI and the safety of electrical devices.

ASIC Application Specific Integrated Circuit.

bit error rate $\frac{\text{Number of errors}}{\text{Number of bits written or read}}$

BOM Beginning Of Media. The first point on the tape that can be accessed by the drive.

buffered mode A mode of data transfer in write operations that 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 that are incorrect.

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 in order to identify errors.

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. *See also* redundancy.

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 LTO-DC 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.

decompression A procedure in which the original data is generated from compressed data.

drop-in Previously recorded data in the midst of new data, which was not been overwritten, probably because of a head-clog.

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.

ECMA European Computer Manufacturers Association. The European equivalent of ANSI.

EEPROM Electrically Erasable Programmable Read-only Memory.

EOD End Of Data. An area that signifies the end of the valid data. If new data is written over a larger quantity of old data, it is possible for data to exist after EOD, but because it is after EOD, this old data is no longer valid.

EOM End Of Media format. The last usable point on the tape.

EW-EOM Early Warning End Of Media. A physical mark or a device-computed position on the tape that tells the drive that it is approaching EOM.

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 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.

hard error An uncorrectable data error.

head clog Particles from the tape or from outside the drive adhering to the head gap on a read or write head that obstruct the reading or writing of data. The particles will often become dislodged again with continued use.

host The host computer system acting as controller for the drive.

load The process in which the drive takes in an inserted cartridge and goes online.

LUN Logical Unit Number, by which different devices at a particular SCSI ID can be addressed individually. The drive has a fixed LUN of 0.

LVD Low-Voltage Differential. *See* SCSI.

noise Any kind of unwanted magnetic or electric interference detected by the electronics.

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 that would cause tape motion. The host can, however, load a tape, if one is inserted, and can execute any diagnostic tests that do not require tape motion.

online The drive is online when a tape is loaded. The host has access to all command operations, including those that access the tape, set configurations and run diagnostic tests.

randomizing A recoding of data symbols before they are written to tape in order to provide a consistently high RF envelope level. An inconsistent 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 also* bit error rate.

read-after-write RAW improves data integrity by reading data immediately after it is written and writing the data again if an error is found.

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.

RF envelope A waveform composed of the instantaneous peak values of an alternating signal that indicates the variation in peak amplitude of the signal.

SCSI Small Computer System Interface—a standard command specification and command set that enables computers and peripherals to communicate with each other. HP's Ultrium drives adhere to the SCSI specifications (see Chapter 1, "Interface Implementation" in Volume 3, The SCSI Interface, of the HP Ultrium Technical Reference Manual) and support all features required by those standard.

Single-Ended and Low Voltage Differential SCSI

These terms define how the signals are transmitted along the cable.

With *single-ended* (SE) SCSI, each signal travels over a single wire and each signal's value is determined by comparing the signal to a paired ground wire. Signal quality tends to decrease over longer cable lengths or at increased signal speed.

With *low voltage differential* (LVD) signaling, signals travel along two wires and the difference in voltage between the wire pairs determines the signal value. This enables faster data rates and longer cabling with less susceptibility to noise than SE signaling and reduced power consumption.

Narrow and Wide, Fast, Ultra and Ultra2 SCSI

Narrow SCSI devices can transfer data one byte at-a-time (and are sometimes called "8-bit SCSI" devices). They can conform to either the SCSI-2 or SCSI-3 protocols. They have a 50-pin connection to the SCSI bus.

Wide SCSI devices can transfer two bytes of data simultaneously ("16-bit SCSI"). They usually have a single, 68-pin connection to the SCSI bus. (This physical arrangement is part of the SCSI-3 specification.) They may support either SCSI-2 or SCSI-3 protocols. Wide and narrow devices can simultaneously be connected to the same bus without problem, provided certain rules are followed.

Fast SCSI can transfer data at up to 20 MB/s wide, using a cable of up to 6 meters total length.

Ultra SCSI can transfer data at up to 40 MB/s wide, but the cable length cannot exceed 3 meters (it is also known as "Fast20").

Ultra2 SCSI can transfer data at up to 80 MB/s wide, using a cable of up to 25 meters total length for a single device, or up to 12 meters for two or more devices (it is also known as "Fast40").

Ultra3 or *Ultra160* can transfer data at up to 160 MB/s wide. Cable lengths are as for Ultra2.

Ultra4 or *Ultra320* will transfer at up to 320 MB/s.

Ultra SCSI supports both SE and LVD interfaces. In normal situations, slower devices can coexist with faster devices, and narrow devices can be used on the same SCSI bus as wide devices using a suitable adapter.

HP's Generation 1 Ultrium drives are Ultra2, wide SCSI-3 compatible devices. They can be used with both LVD and SE host bus adapters.

Generation 2 Ultrium drives are Ultra160, wide SCSI-3 compatible.

single-ended *see* SCSI

soft error A soft error is a data error that can be corrected by a RAW rewrite during writing to tape, or by a read-retry during reading.

spacing Spacing is moving along the tape over a specified number of blocks or filemarks, or to EOD, in order to find data 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).

tape mark A filemark.

TapeAlert A set of 64 flags is held in the TapeAlert log that indicate faults or predicted faults with the drive or the media. By reading this log, host software can inform the user of existing or impending conditions, and can, for example, advise the user to change the tape.

vendor-unique The addition of commands to SCSI that are not included in the standard.

index

Index