

HP StorageWorks Ultrium 960 tape drive technical white paper



Executive summary.....	2
LTO Ultrium roadmap.....	3
HP StorageWorks Ultrium 960 tape drive—new features.....	4
The LTO Ultrium 3 format.....	5
HP StorageWorks Ultrium 960 controller electronics.....	10
SCSI-3 parallel interface.....	11
Fibre Channel interface.....	12
ACI.....	12
HP StorageWorks Ultrium 960 PRML read channel.....	13
ALDC data compression.....	15
Data rate matching.....	15
WORM tape.....	16
Why WORM tape?.....	16
How does it work?.....	16
HP Ultrium 3 media and compatibility.....	18
Conclusions.....	19
References.....	20
For more information.....	20

Executive summary

The Linear Tape Open (LTO) Ultrium format was introduced in 2000 by HP, IBM, and Seagate and has become established as the de facto standard in the “superdrive” marketplace. A comprehensive roadmap was published at the launch of the format, and the HP StorageWorks Ultrium 960 tape drive introduces the third generation of LTO Ultrium technology.

With a native capacity of 400 GB and a throughput of 80 MB/s, the HP StorageWorks Ultrium 960 represents the highest capacity and performance solution of the LTO Ultrium family. Designed for mid-range and enterprise-level customers who need this standard of performance and reliability for direct attach data protection solutions, the LTO Ultrium 3 format remains an open standard, ensuring interchange of tape media between any product that carries the Ultrium compliance mark.

Award of the Ultrium logo means that drives or tape media have passed the strict compliance verification program. The HP StorageWorks Ultrium 960 tape drives are fully read compatible with Ultrium 1 products and read-write compatible with Ultrium 2 products. This ensures that customers’ investment in previous Ultrium technology is protected and that archive tapes can still be accessed. The new write-once read-many (WORM) technology has been introduced to the HP StorageWorks Ultrium 960 tape drives. Using WORM data cartridges customers can ensure that data cannot be overwritten, erased, or altered. This standard feature provides HP customers with a simple, secure, and cost-effective method for archiving data where legal compliance is required. HP has incorporated many technical design improvements to both the drive mechanism and the electronics module of the Ultrium 960. These improvements, based on the experience of the first two generations of Ultrium technology, take reliability and performance to a new high point for stand-alone or automation applications. Customers choosing HP LTO Ultrium 3 technology will find it more than matches the ever-increasing demands of today’s data protection solutions.

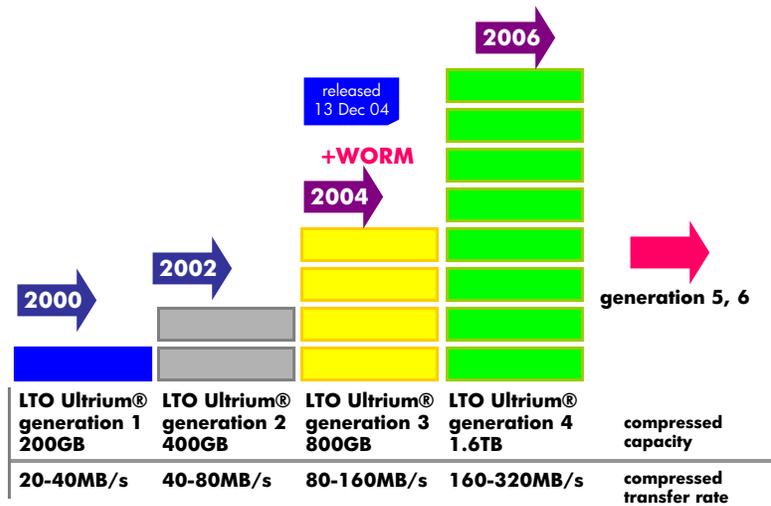
Figure 1. HP StorageWorks Ultrium 960 tape drive—desktop version



LTO Ultrium roadmap

The LTO Ultrium format has a well-defined roadmap (Figure 2) and extends well into the future with at least a doubling of capacity and throughput for each generation. Maintaining backwards compatibility is a design goal to ensure investment protection in tape media and to enable archive material to be read at later dates on future generation Ultrium drives should this be necessary.

Figure 2. LTO Ultrium roadmap



HP StorageWorks Ultrium 960 tape drive—new features

The HP StorageWorks Ultrium 960 tape drive introduces the LTO Ultrium 3 format for the first time. The technical specification of this format is laid down by the ECMA U-316 document. This format delivers a native capacity on a single tape cartridge of 400 GB. Throughput has been increased to 80 MB/s. Full read-write compatibility with Ultrium 2 is provided with read-only compatibility for Ultrium 1 format tape media. Media format is selected automatically when a data cartridge is inserted. It is not possible to write different formats on the same media. A new yellow cartridge makes it easy for users to identify HP Ultrium 3 tape cartridges. The cartridge color-coding scheme used by HP is not universal and some brands of media are either black or different colors.

To cope with the increased throughput the HP StorageWorks Ultrium 960 tape drives are equipped with an Ultra 320 parallel SCSI interface. A dual 2-Gbps Fibre Channel interface will only be available on the tape library drive mechanisms that are integrated into the HP StorageWorks tape library family. Only Ultra 320 SCSI can deliver the throughput required by a tape drive of this performance. Data is compressed by the drive hardware so typical 2:1 data would require a bus transfer rate of 160 MB/s. When overheads are taken into consideration, lower speed SCSI busses would not deliver the full performance. A new larger SCSI burst buffer and main data buffer help balance the data flows to and from the server.

Several new features have been necessary for the tape drive mechanism to handle the increased demands of the format. New profile tape rollers and a tape lifter mechanism are designed to improve reliability. The tape lifter mechanism lifts the tape leader away from the read-write head surface during the cartridge load and unload cycle. This prevents wear on the head tape-bearing surfaces and prolongs head life. The tape head cleaner is unchanged and works in conjunction with the universal intelligent cleaning cartridge. A vibration sensor is now attached to the drive mechanism and its output used by the servo electronics to increase the accurate head positioning over the data tracks.

To accommodate the Ultrium 3 format and to maintain backwards compatibility a new 16-channel head design is used. The head uses magneto-resistive (MR) read heads and thin film inductive write heads. Four MR servo heads are used to locate the data heads on the designated tracks. The head linear motor assembly and variable inductance sensor are largely unchanged from previous HP Ultrium 1 and Ultrium 2 tape drives.

The Data Rate Matching range is changed in the Ultrium 3 format and can vary the tape speed in a range to produce a native data rate of between 27 and 80 MB/s.

Completely new to the Ultrium 960 tape drive is WORM media. This uses servo code and cartridge memory (CM) features in a special cartridge to implement a WORM facility that prevents data from being erased or overwritten and can be used as part of a legally compliant data protection solution. WORM is logo certified and can be interchanged with other vendors' WORM Ultrium 3 tape drives and vice versa.

Table 1: HP StorageWorks Ultrium 960 tape drive—new features

New feature	Benefits
Ultrium 3 format tape drive with full read-write compatibility with Ultrium 2 and read compatibility with Ultrium 1. Ultrium 3 is the third generation of the LTO Ultrium standard	400-GB native capacity that is double that of the previous Ultrium 2 format with much greater throughput (80 MB/s)
Ultra 320 SCSI interface with dedicated and more powerful microprocessor	Enables faster data transfers between host systems and the tape drive, which can dramatically reduce backup and restore times
16-channel head design	Required for the Ultrium 3 format
Increased main buffer size to 128 MB	Enables better data flow and efficiency particularly if the host server has varied performance
SCSI burst buffer increased to 4 MB—actually 8 MB but only 4 MB is used	Essential for efficient handling Ultra 320 data bursts
Enhanced Data Rate Matching range—27 MB/s–80 MB/s	Improved performance with slow host servers
Vibration sensor (accelerometer) attached to the drive mechanism	Improves the tracking accuracy of the head servo mechanism and increases the tape drives immunity to shock and vibration
Mechanism improvements including a tape lifter device, improved roller coating, and alignment	Tape lifter minimizes head wear during load/unload cycles Higher tape drive reliability
Improved drive electronics module with less subassemblies	Greater reliability and performance
Enhanced automation control interface (ACI) with negotiated baud rate to 153.6 Kbaud	Enables automation applications to enhance functionality
WORM tape media capability Special two-color easily identifiable data cartridge	Ideal for enhanced data security and legal compliance
Enhanced built-in diagnostics and monitoring capability	Enables greater functionality in diagnostic utilities such as HP Library and Tape tools
New data compression core	Increased data throughput

The LTO Ultrium 3 format

In discussing the tape format, it is important to note that this is an open standard and therefore applies to all vendors. Tape drive performance is independent of format, and vendors can set their own specifications. Certified Ultrium products will carry the Ultrium logo and are independently verified for format compliance.

The new LTO Ultrium 3 format is similar to the previous Ultrium 1 and 2 formats and consists of serpentine linear recording with a timing-based pre-recorded servo code. The physical outside dimensions of the cartridge are identical to maintain media compatibility, but the tape length has been increased to 680 m ± 1 m. The tape width remains the same at 12.65 mm ± 0.006 mm. The native capacity has been increased to 400 GB, which has been achieved by increasing the number of data tracks to 704 with a track pitch of 14.3 μm. (Ultrium 2 format has 512 tracks [track pitch

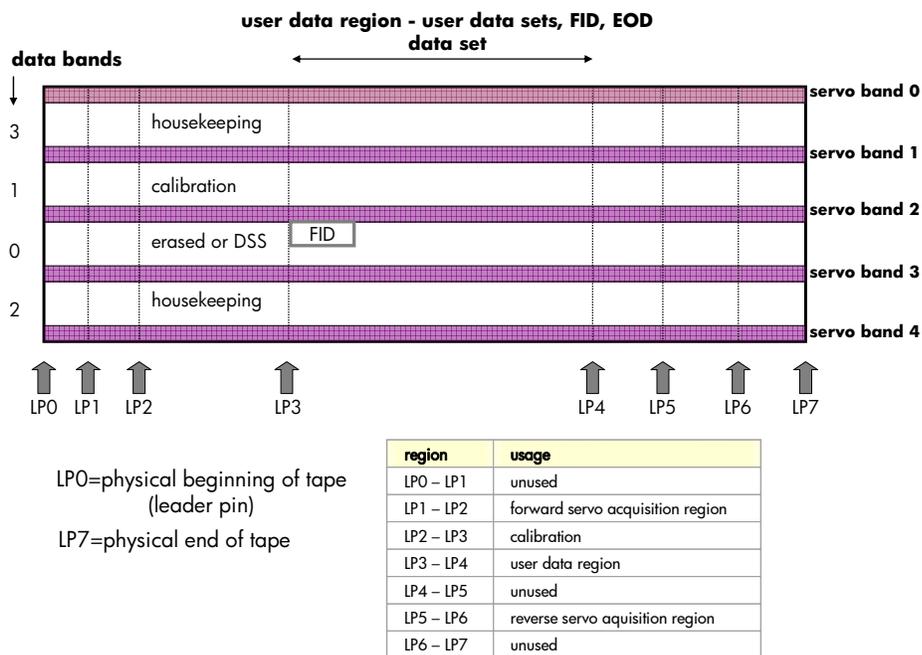
20.17 μm] and Ultrium 1 has 384 tracks). The linear recording density is increased to 243 kilo-bits per inch (kbpi).

Ultrium format tapes are identified by the CM before drive threading the tape. The Ultrium 3 pre-recorded servo code is essentially unchanged other than encoding for the WORM option. Servo code is pre-recorded at the tape media manufacturer and it cannot be altered or erased by the tape drive. Bulk erasure would destroy the servo code and render the cartridge unusable. The timing-based servo code was designed from the outset to cope with higher data densities and also to enable backwards compatibility. The CM is unchanged in capacity but does contain specific format data that is not accessible by the user. CM is an integral part of the Ultrium format.

The Ultrium 3 format uses a write equalized (0,13/11) run length limited (RLL) code. This is unchanged from Ultrium 2 but is different to Ultrium 1 format, which used an (1,8) RLL code for recording data. Extensive Reed Solomon block code error correction is used to maintain data integrity. Ultrium 3 format retains four data bands between a pair of servo tracks, but the band now contains 176 data tracks. A 16-channel read-write head is used in Ultrium 3 format, which has changed from the eight-channel heads used in Ultrium 1 and 2 formats. All read-write data channels are "read after write" to ensure data integrity and to ensure that defective areas of tape to be rewritten or in the worst case a read or write error reported. The uncorrected error rate for the Ultrium format is 1 in 10^{17} bits when reading data.

The tape is divided linearly into seven regions, which are marked by eight logical points. Figure 3 shows the location of these logical points.

Figure 3. Ultrium format tape layout



When a tape is loaded, the CM is read and if the tape is compatible, it will load to the start of the user data region. The use of the calibration area is not defined in the format. It may be used by a tape mechanism vendor to set up parameters such as write currents. As the tape crosses LP1, it will lock on the servo tracks in preparation for aligning the heads over the correct data track area. The format identification data set (FID) is written on initial cartridge load and can be used by the drive if the CM is found to be defective. Note that with Ultrium 3 WORM, there is a slightly different sequence of events and these are described in the WORM tape section of this white paper. The tape length is stored in the CM and when known, the differential reel speeds determines the position of the heads along the tape.

The servo system in the Ultrium 3 format is essentially unchanged and as previously stated is capable of handling LTO Ultrium 1, 2, and 3 formats as well as the new Ultrium 3 WORM media. Known as a timing-based servo system, it is key to the format and provides constant feedback on the speed of the tape and the position of the data heads in relation to the data tracks. An embedded code system within the servo code provides a constant set of position values known as Linear Position (LPOS). The servo bands are recorded by the tape media manufacturer. Figure 4 shows where the servo bands are located and a graphical representation of the data head. The read, write, and servo head assembly is mounted on a moving carriage with the electrical signals carried to it by flexible printed circuits. The carriage is driven by a linear actuator or motor. When no servo signal is present, a variable inductance sensor provides positional feedback to keep the heads on track. The actuator mechanism in an HP StorageWorks Ultrium 960 tape drive is shown in Figure 5. It is important to note that different Ultrium tape drive vendors can use any form of actuator, but it must have the ability to track the standard format of Ultrium tapes.

Figure 4. Ultrium 3 format format details

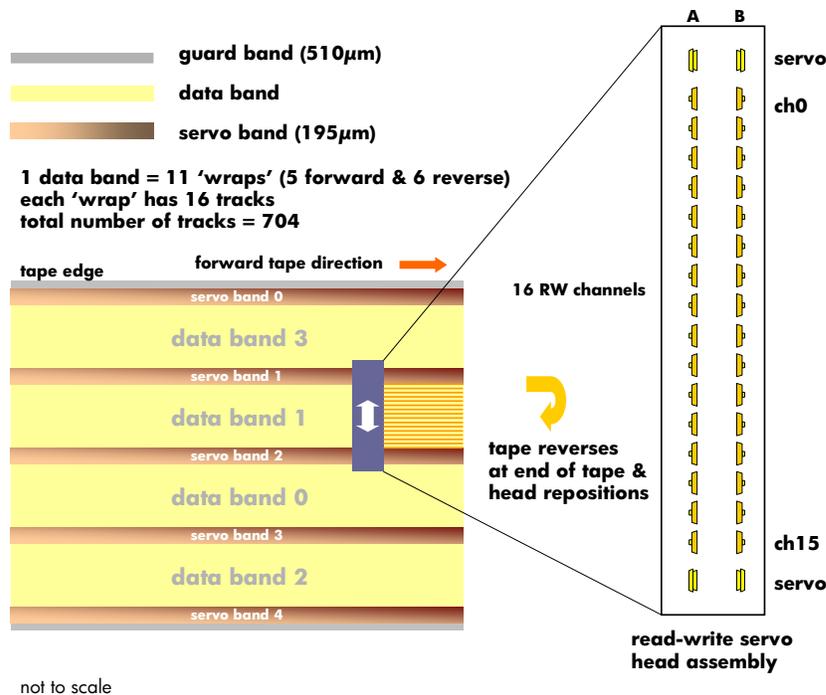


Figure 5. HP StorageWorks Ultrium 960 actuator assembly



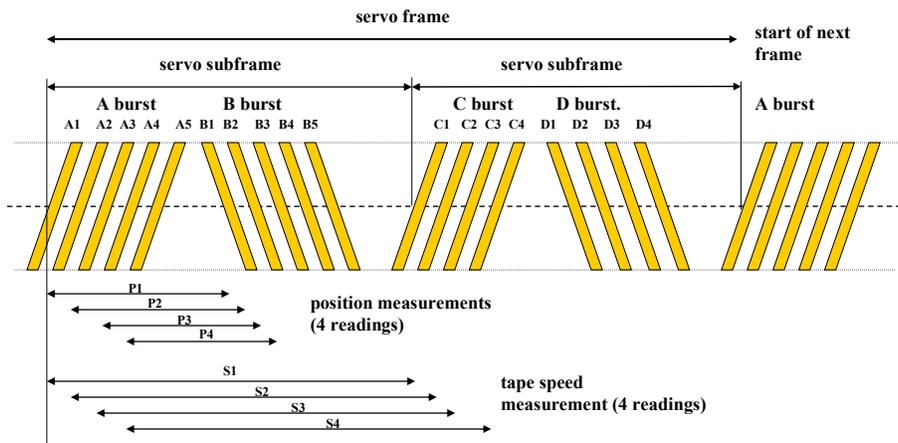
Each servo band consists of magnetic stripes arranged as in Figure 6. The servo bands are much wider than the head element. They are arranged in frames of 18 stripes in sets of five and four in a 5544 sequence. This sequence is used to aid identification of the servo frames.

The servo tracks provide information from which the head position across the tape path, tape speed, and longitudinal position is derived. The head position is derived by timing the pulses A1-B1, A2-B2, A3-B3, and A4-B4. This equates to four positions per servo frame. The mean time from these signals is used in conjunction with the signal from the servo head, which is located on the other servo band. There are always **two** servo heads positioned on servo bands above and below the data tracks. Four timing signals—A1-C1, A2-C2, A3-C3, and A4-C4—are used to measure the tape speed. In Ultrium 3, 11 head positions are derived from each pair of servo bands. Fewer positions are required in the Ultrium 3 format because of the increased number (16) of data tracks being written simultaneously on each pass. (Each complete pass in a particular direction is known as a “wrap.”) The odd number of wraps is used to maximize the tape width and each servo band alternates five forward and six reverse or six forward and five reverse wraps to preclude the need for a redundant pass.

Track trimming is used in Ultrium 2 and 3 formats, which means that each subsequent set of tracks slightly overlaps the adjacent tracks. This overlapping ensures that no space is lost “between” adjacent tracks. If a servo signal is unreadable during write operations, then the drive will post a check condition and an error status, but if the operation is read only, then it continues operations. To identify servo bands, the servo bursts are offset from each other. Having one servo head positioned in each band enables the relative offset of the signals to identify which data band the head is positioned over.

The servo stripes are also used to define the linear position (LPOS). This is achieved by shifting the position of the second and fourth stripe in the A and B servo subframes. Because both the A and B subframes do this, the relative timings are unchanged so that the head position and speed measurements are unaffected. A complete LPOS word is derived from 36 servo frames. The LPOS information is pre-written when the servo tracks are written during tape manufacture. The makeup of an LPOS word is shown in Figure 6. The longitudinal position is just a number that increases steadily along the track. When a tape cartridge is first loaded, the LPOS information is written to the CM. This enables BOT to be located in each subsequent cartridge load. The LPOS number is also used to encode information such as tape vendor ID within the servo information. Because the LPOS information is known to the drive, "drop-ins" are prevented. Drop-ins occur when there is a head clog on a write channel and old data is left on the tape. In Ultrium tape drives, this is immediately identified and the appropriate action taken. The new Ultrium 3 WORM format also uses the information encoded within the LPOS signal to identify WORM media (see the WORM tape section). The HP StorageWorks Ultrium 960 tape drives are compatible with Ultrium 1, 2, and 3 formats (Ultrium 1—read only) and the servo electronics is programmed appropriately after the tape is identified from the CM (or FID in the case of CM failure). Tapes with a defective CM can be read to recover the data but writes are not allowed.

Figure 6. Ultrium format—servo band detail

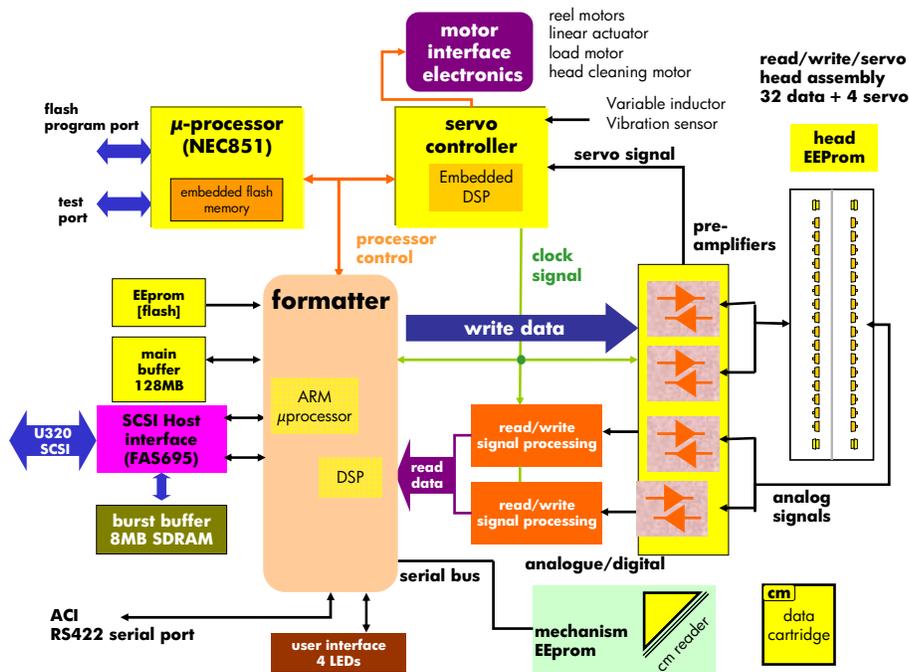


HP StorageWorks Ultrium 960 controller electronics

The new HP Ultrium 3 tape drive electronics is based on a customized 32-bit NEC microprocessor core with totally new ASIC "front-end" formatting and signal processing sections. New interface electronics with a separate microprocessor provide either Ultra 320 SCSI or 2-Gbps Fibre Channel interfaces. A new design for the automation control interface (ACI) has been implemented to improve performance and functionality. The electronics are mainly contained on a single printed circuit assembly located beneath the tape drive mechanism. Drive airflow has been designed to flow past this PCA but not through the drive mechanics. This directs cooling air to where it is needed and keeps the drive mechanism free from dust and airborne contamination. Firmware code is held in flash memory and can be updated by the host interface, by the ACI, or from a specially formatted tape. A block diagram of the controller electronics is shown in Figure 7.

The front-end electronics of the HP Ultrium 3 tape drive is a totally new design compared to the HP Ultrium 1 and Ultrium 2 tape drives. The HP Ultrium 3 tape drive has a 16-channel read-write head design with 32 MR read-write elements and four MR servo elements. Increased bandwidth is required as the linear bit rate is now 243 kbp. The new design of preamplifier can cope with a bandwidth of 100 MHz and has individual bias control for each MR read element. As MR heads are based on a semiconductor junction, they require a voltage bias to optimize the output. Each pre-amplifier section can handle eight channels, therefore, four sections are required to cover all the heads. The signals from the servo heads use separate four channel selectable pre-amps. Servo inputs are processed and fed to the linear actuator and the tape reel motors. The trailing servo heads are used so as to prevent crosstalk from the write elements. Tape speed information is derived from the servo code but also from the differential speeds of the tape spools.

Figure 7. HP StorageWorks Ultrium 960 block diagram



A chassis-mounted accelerometer is used as a vibration sensor and is used to provide feedback into the servo system to improve the ability of the head assembly to stay on track. The servo controller also drives the load/unload motor and the head cleaning motor.

SCSI-3 parallel interface

The HP StorageWorks Ultrium 960 tape drive has a parallel SCSI interface available for internal, desktop, and rack mount versions. The interface is capable of transferring data bursts at 320 MB/s. This is commonly called Ultra 320 SCSI. The SCSI-3 parallel interface is 16 bits wide only.

The SCSI standard for the interface complies with SAM-2, SPC-3, and SSC-2 standards. These are the standards of the T10 committee of International Committee on Information Standards (INCITS) and operate under the North American ANSI rules (<http://www.t10.org>).

The HP StorageWorks Ultrium 960 parallel SCSI interface hardware is based around the QLogic FAS695 fast architecture SCSI processor ASIC. The FAS695 SCSI controller has built-in low-voltage differential and single-ended transceivers. HP has used an external ARM9 microcontroller running at 110 MHz. The SCSI interface is designed to handle bursts of data up to 320 MB/s. The FAS695 ASIC has a high speed "bursting" DMA buffer controller that is used to interface to a 64-MB SDRAM burst buffer. This is separate from the main buffer of the tape drive, which is 128 MB. The burst buffer can handle the extremely fast 320-MB/s bursts of data from the SCSI bus. Data can then be transferred to the drive electronics at a more constant rate.

Ultra 160 SCSI saw the introduction of dual edge clocking, domain validation, and cyclic redundancy check (CRC). Ultra 320 SCSI has several new additional protocol features that enhance performance. These changes improve performance and reduce overhead, allowing safe data transfers at 320 MB/s. New features of the SCSI interface that are implemented in the HP StorageWorks Ultrium 960 tape drives include:

- Double transfer speed to 320 MB/s
- Target initiated negotiation—This is useful when hot-plugging devices, power cycling, or moving devices between separate SCSI busses. It enables re-negotiation to be initiated by the device and not exclusively the host bus adapter (HBA).
- Packetised SCSI—This is mandatory. Packetised devices decrease command overhead by transferring commands, data, and status using the dual transition data phases instead of slower asynchronous phases.
- Tagged queuing—Is supported but with a queue depth set at 1. This feature is required by certain vendor's HBAs (for example, Adaptec).
- Pre-compensation—This new signaling technology is required to maintain high reliability at the higher bus speeds. Although signals at Ultra 320 speed are twice the frequency of those at Ultra 160 speeds, cable requirements are essentially unchanged as a result of the pre-compensation of the signal in the bus transceivers. (However, HP strongly recommends the use of specifically designed SCSI Ultra 320 internal ribbon cables.)

The following SCSI features are **not** supported in the HP StorageWorks Ultrium 960 tape drives:

- SCSI configured automatically (SCAM)
- Quick arbitration
- Read and write data streaming

For maximum performance, it is recommended that Ultra 320 HBAs are used. The data transfer rate of the HP StorageWorks Ultrium 3 tape drive is 80 MB/s without data compression.

Fibre Channel interface

A class 3 Fibre Channel version of the HP Ultrium 3 drive has been developed for OEM marketplace and is used in tape automation applications. The HP Ultrium 3 tape drive uses a second-generation QLogic FAS495 dual channel ASIC with a more powerful separate ARM microprocessor. The jumpers located next to interface no longer set the AL-PA addressing. The hard address can be set by way of the ACI port. However, arbitrated loop connections are rarely used today. The Fibre Channel interface manager implements the FC AL-2 protocol and does of course support fabric login. Dual LC Fibre Channel connectors are provided on the rear of the drive. For automation vendors, it is possible to download a new worldwide name (WWN).

ACI

In addition to the standard interface, an RS422 serial interface is provided for automation control. This enables tape libraries and autoloaders to communicate with the tape drive independently from the parallel SCSI or Fibre Channel interface. The ACI port in the HP Ultrium 3 drive can now support a faster baud rate and is bi-directional with XON-XOFF flow control to 115,200 baud. The ACI is also capable of accepting surrogate SCSI commands. This feature is useful in the autoloader application when only one SCSI interface is required. The ACI in the HP Ultrium 3 drive is much faster than the previous generations and provides immediate acknowledgment of transmission. The ACI port uses a PH-9 connector.

The data flows through the controller as follows:

1. Writing data to tape.

Data is received from the host interface by way of Fibre Channel or Ultra-320 parallel SCSI. The data arrives in fast bursts of data depending on the setup of the host software, and is held in the burst buffer (SSRAM). The burst buffer is now 4 MB and is external to the ASIC buffer. The main purpose of this buffer is to isolate the high-speed burst nature of the SCSI data from the rest of the controller electronics. It should not be confused with the main drive buffer, which is 128 MB but is not located in the interface electronics. Data is taken from the burst buffer by the formatter and has record boundaries added. Data then has a CRC record added to ensure that no errors remain undetected by the formatter process. The data is compressed by the ALDC core as described earlier and has a C1 parity data added before being stored in the main buffer. This ensures that errors in the main buffer memory are detected in the case of data corruption by the SDRAM. The purpose of the main buffer is to provide a steady flow of data to the recording electronics. The data is then taken from the main buffer and has powerful C2 Reed Solomon codes added. Before passing the data to the CCQ (codeword quads) writer, the C1 error detection picks up any errors that may have occurred in the main buffer SDRAM. The CCQ writer then randomizes that data to remove long sequences of ones or zeros, and then encodes the data using an RLL code. Because of the higher linear density, a different RLL 0,11/13 code is used for the Ultrium 3 format. The signals are then subject to write equalization before being applied to the write head elements by way of the write pre-amplifier. Remember 16 channels are written simultaneously. The data is immediately read by the 16 read elements on the trailing head "bump" and sent back by way of the digital signal processor to the physical formatter. If the C2 code is in error, then the data block is rewritten on the tape. This is known as a read-while write retry. The data block rewrites are always on a different channel so that small media defects are easily handled.

2. Reading data from tape.

The read process is basically a reverse of the write process. The magnetic signal is retrieved from the tape by the 16 MR elements of the tape head. (The servo code is always being read and fed into the servo system.) The analog data is amplified and filtered before being fed into the digital signal processor ASIC. The signal is then sampled at intervals and an analog-to-digital converter is used to convert the waveform into sets of 8-bit numbers. The Viterbi detector uses maximum likelihood techniques to recover the data. The data is passed to the formatter, which decompresses and restores the data to its original form and sends it to the interface electronics.

HP StorageWorks Ultrium 960 PRML read channel

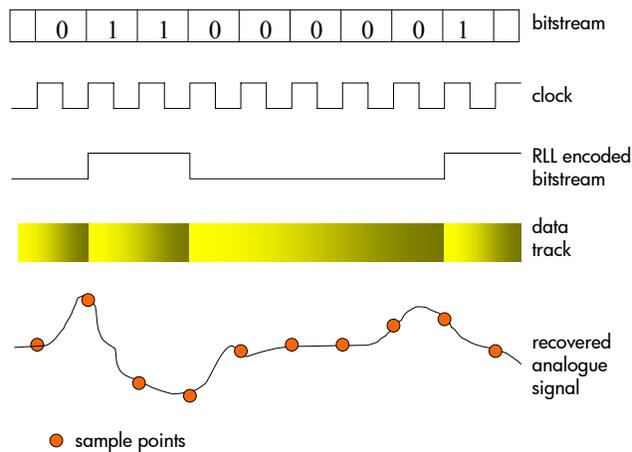
The read channel of the HP StorageWorks Ultrium 960 uses Partial Response Maximum Likelihood (PRML) technology to recover the data signal from the tape. This well-established technology is essential to recover data at the densities used in the Ultrium 2 and 3 formats. PRML has been used to increase the capacity of disk drives for many years. The HP StorageWorks Ultrium 230 tape drive uses a peak detect method of recovering the digital data from the analog signal and, although well capable at this, data density is not sufficient for the densities of data used in Ultrium 2 and 3 formats. To describe the read channel it is necessary to understand how the data is encoded before it is written to tape. After data compression and adding the Reed Solomon error correcting coding, the data is in the form of CCQ. Reed Solomon encoding is a well-known block code, which has excellent error recovery methods. The data is then run through a data scrambler, which removes long strings of zeros. These are undesirable when writing magnetic information, as the read head output is proportional to the rate of change of magnetic flux. The data is then encoded using an RLL code before equalization and then turned into a write current. Ultrium 1 format used an RLL 1,7 code but this has changed in Ultrium 2 and 3 to an RLL 0,13/11 code. This will encode no more than 13 consecutive bits as a zero before changing state and no more than 11 consecutive interleave bits.

This is more suitable to magnetic flux changes at double the density of the Ultrium 1 format. It is important to note that eight channels are encoded simultaneously as the Ultrium 2 drive uses an eight-channel head. In Ultrium 3 format 16 channels are encoded simultaneously.

The bit stream is now written on the data tracks by the write heads. The read head recovers the signal from the tape but is subject to electrical noise from the media. As well as noise, the density of the flux changes in relationship to the data head "gap," causing inter-symbol interference that peak detect systems cannot handle. The signal is amplified and subject to equalization so that the frequency response is more even. The data is written with a "clock" signal, which is recovered from synchronization bursts occurring at intervals in between data signals. A phase lock loop system is used to maintain this signal. The clock signal informs the read decoder the position of the bit cell transitions. The signals are shown in Figure 8 with a typical bit stream. The signal is then sampled at the intervals provided by the phase locked loop and is turned into a 36-bit digital number. This is the partial response part of the process. The signals are then digitally filtered. An advanced Viterbi detector is used to maintain a history (36 bit long, 16 states wide). The Viterbi detector makes decisions on the data sequences received, based on the history it has built up. It operates similarly to a person reading a document with spelling mistakes, but still being able to accurately understand the text. Any errors would of course be corrected by the powerful C1 and C2 Reed Solomon error correction, which is used after decoding.

After the Viterbi detector, the digital bit stream is decoded from the RLL code and unscrambled. The signals can then be returned to the formatter and after de-compression returned as blocks of real data to the I/O system. The advantage of this system is that it can also recover Ultrium 1 format signals as the digital filtering can be re-programmed and the decoding changed to cater for RLL 1,7. The drive identifies an Ultrium 1 tape from the CM system.

Figure 8. HP Ultrium 3—PRML data sampling



ALDC data compression

The data compression method used in Ultrium technology is defined in the format. Therefore, all vendors must use the Advanced Lossless Data Compression (ALDC) algorithm for data compression. There is no change in the Ultrium 3 format standard. ALDC is an implementation of the Liff Zempel method of compressing data and can be implemented in hardware more efficiently than the standard DCLZ method of data compression. In itself ALDC does not compress data any more efficiently than the DCLZ method, but has the ability to switch into a non-compressed mode according to the structure of the data pattern. This means that highly random data does not actually expand when compressed. ALDC in the HP StorageWorks Ultrium 960 tape drive is implemented in hardware within the formatter ASIC before data is sent to the main I/O data buffer. The HP Ultrium 1 and HP Ultrium 2 tape drives used a contents addressable memory (CAM)-based data compression engine. The HP StorageWorks Ultrium 3 data compression is synthesized within the formatter ASIC. Compression and decompression are totally separated and with the increased clock speeds used has a maximum data rate of 240 MB/s. This means that the drive can handle a constant stream of data that can be compressed at a 3:1 ratio.

The stream of data is received from the SCSI interface electronics and read into a 1-K buffer. The data by this time is in the form of records. The data is read into a shift register. The data is output from the buffer as either "literals" or copy pointers. The literals are 9 bits in length, with the most significant bit set to 0. The copy pointers have the most significant bit set to 0 and contain a length variable and an address within the shift register. The shift register is 1 K x 9-bit words. At record boundaries the process is reset and started over again. The compression rates are monitored and if compression is not being achieved, the drive switches modes and only literals are put into the data stream. When data is found to be compressible, the compression is turned back on.

Data rate matching

The HP StorageWorks Ultrium 960 tape drives use a feature known as data rate matching (sometimes referred to as adaptive tape speed [ATS]) that can dynamically vary the speed of the tape according to the data flow rate to or from the host system. This technology was successfully used in the Ultrium 1 and 2 tape drives and has now been enhanced in the HP StorageWorks Ultrium 960 tape drive.

The objective of the adaptive speed algorithm (ATS) is to minimize the number of repositions performed by the tape mechanism. A reposition happens when there is either insufficient data from the host or when the host cannot accept the data from the drive. After stopping, the tape drive must reverse direction, go back beyond where the last data was, and then go in the forward direction again. Repositioning the tape wastes time and causes extra wear on the media and tape mechanism. ATS continually monitors the host data rate and calculates a "set point" within the main data buffer at which the speed of the tape is varied to maintain the level of the data within the buffer. The set point position varies according to the data rate and when the end of wrap approaches. At the end of wrap during write, the drive will ensure the buffer is nearly empty if possible so the buffer space holds data while the tape turnaround is executed. During a restore, the ATS system will keep the buffer nearly full before the end of wrap so that data flows out of the buffer back to the host. The main buffer on the HP StorageWorks Ultrium 960 tape drive is 128 MB. (HP StorageWorks Ultrium 1 drives are 16 MB.) The ATS range for the Ultrium 960 tape drive is 27–80 MB/s for Ultrium 3 tapes. The differences in ATS ranges are because of the different data throughput for each format. (Ultrium 3 has greater tape speed and data density.) The range of the ATS feature is dictated by read-write channel's ability to cope with the bandwidth of the read-write channel. (The read signal below a certain speed would be too low and at high speeds the frequency starts to hit limits.)

Of course, if the host data rate is out of the ATS range, the buffer will be used until eventually the tape must stop and a reposition is necessary. The clever feature is that ATS is dynamically adaptive and does not stay at a fixed buffer level. ATS makes the HP StorageWorks Ultrium 460 ideally suited for backup over a network where data rates can vary.

WORM tape

The term WORM first appeared with magneto-optical disks in the 1980s. Disks were 5.25-inch or 12-inch media and used a laser in conjunction with a magnetic field to write data and a laser to read the data. WORM media magneto-optical disks are pre-formatted in a manner that prevents overwrite after the data has been written. Although a non-ablative technology, magneto-optical WORM has been legally compliant in most countries as media on which data, after written, cannot be modified. Magneto-optical disks have now been superseded by Ultra Density Optical (UDO) disk, which although in a similar package to magneto-optical media, use a phase change technology and laser only. In WORM versions of UDO media, the change to the media is permanent and cannot be reversed.

Why WORM tape?

Ultrium tape technology has much greater capacity than even the latest UDO technology. It has an archive life of 30 years and is a proven stable media. The current writable CD and DVD technology has a much shorter archive life than this and is prone to data errors. Long-term storage of CD and DVD media is certainly unproven. Tapes that have been stored in the correct conditions have been shown to be stable and readable after many years. Increasingly stringent legal compliance has made WORM-type storage a vital part of data protection.

WORM tape technology is a standard feature with the HP StorageWorks Ultrium 960 tape Ultrium 3 tape drives. WORM Ultrium tape media from HP has a distinctive two-tone yellow and grey color and when used will only permit writing once to the media. Users can of course append data on the unused section of the tape until the tape is full.

How does it work?

The HP WORM tape technology uses a multi-level security system and a special WORM tape cartridge that prevents any overwrites of existing data. The distinctive two-color WORM data cartridge is physically identical to standard read-write media cartridge but has a unique cartridge-type code that is written into the CM at manufacture. This is a read-only field so it cannot be modified. If loaded in a non-WORM-compatible tape drive, the cartridge will be ejected and the tape error LED will flash. In tape library applications, the cartridge would not actually eject since auto-eject is normally disabled. In these applications, the cartridge is held in the "eject ready" position. The cartridge type also defines the LTO generation with which the media is compatible.

The CM is uniquely tied to the data on tape as a security feature against tampering with the cartridge (for example, removing the tape spool and playing in a normal read-write cartridge). This is accomplished by copying the Cartridge Manufacturer's Page, the Media Manufacturer's Page, and the Initialization Data Page read from the CM into the Format Identification Data Set (the first data set written to tape). As a result, the data on tape contains information that unambiguously indicates that the CM and the tape in the cartridge are associated. Specifically, these pages contain information such as the tape batch identification from the Manufacturer's Word in the Servo track and LPOS values that mark specific areas on tape. (This feature is already within the LTO format.)

When a cartridge is loaded, a Unique Cartridge Identity (UCI) word is formed that consists of the following fields: LTO CM Serial Number, Pancake ID from the Servo code, and the LPOS value at LP1. This UCI is used to uniquely identify a WORM cartridge and will be incorporated within the Data Set Information Table (DSIT) of every data set written to tape. In addition, the UCI will be accessible to software applications by way of the MAM read attribute command.

MAM attribute 1000h will return the UCI on an initialized cartridge. The length is 28 bytes.

Important:

As the UCI includes the LPOS value and that this value is not finalized until the cartridge is initialized, then the UCI will not be returned for an uninitialized cartridge and the command will be check conditioned with sense 03/1112h (Medium Error/Auxiliary Memory Read Error).

To provide additional security against tampering with the cartridge, the UCI word will be copied into the DSIT of every data set written to tape to associate the UCI stored within the CM with each data set written to tape.

A new character is added within the servo code of LTO3 cartridges that is used to identify that the media is either of type RW (read-write) or WORM. This character will be placed at a specific location in the Servo Manufacturer's Word and will be used as part of the integrity check for the cartridge.

The Write Protection TAB, or WORM tape cartridges, is retained and behaves normally. When in the write-protect position, any write attempt will produce a check condition.

In normal operational conditions, a WORM cartridge will behave as follows:

1. WORM cartridge is inserted in WORM-compatible drive.
2. CM identifies it as WORM media.
3. Drive reads EOPD from the CM.

The tape drive keeps a record of the End of Protected Data (EOPD) value, which is a logical position on the tape. The tape between the beginning of the media and the EOPD cannot be overwritten. This value is retained in CM and the tape drive updates its copy as each block of data is written to the tape.

The tape drive will accept write (records or filemarks) commands only if the current logical position is beyond the EOPD value. If a write command is received and the logical position is not beyond EOPD, then the command is rejected and a request sense will return data protect/WORM media—overwrite attempted. Tape alert flags and write protect are set.

Erase commands (short or long) to a WORM-compatible drive with WORM media loaded will be rejected and a check condition set. The request sense will return the same status as an overwrite attempt.

If the logical position is beyond EOPD, then write commands are permitted as normal until the end of the tape is reached. A copy of the UCI word is written to tape and is incorporated into the DSIT for every record.

The HP StorageWorks Ultrium 960 tape drive has WORM capability as standard and only requires the use of the easily identifiable WORM media to provide a secure medium for long-term storage, which may need to meet legal compliance standards.

HP Ultrium 3 media and compatibility

HP Ultrium 3 media is color-coded yellow with the WORM variant having the lower half of the cartridge shell in grey. The tape length is extended to 680 m (Ultrium 2 to 609 m). The tape thickness has been marginally reduced to 8 μm and the coercivity of the magnetic layer is increased to 2,500–2,700 Oersteds. The CM is identical in size but is re-programmed with the Ultrium 3 format values and if WORM will indicate WORM media. Table 2 defines the compatibility of the media.

Note that HP and other media vendors sell media labeled at 2x native capacity, that is, Ultrium 960 media is shown as 800 GB.

Table 2. Ultrium media compatibility

Media type Native tape capacity		Ultrium 1 tape drive		Ultrium 2 tape drive		Ultrium 3 tape drive	
		Read	Write	Read	Write	Read	write
Ultrium 1	100 GB	Yes	Yes	Yes	Yes	Yes	No
Ultrium 2	200 GB	No	No	Yes	Yes	Yes	Yes
Ultrium 3 RW	400 GB	No	No	No	No	Yes	Yes
Ultrium 3 WORM	400 GB	No	No	No	No	Yes	Yes*

* WORM tape will not overwrite, only append. It is only accepted by WORM-compatible drives.

HP media product numbers:

- HP StorageWorks Ultrium 1 200-GB data cartridge—C7971A
- HP StorageWorks Ultrium 2 400-GB data cartridge—C7972A
- HP StorageWorks Ultrium 3 800-GB data cartridge—C7973A
- HP StorageWorks Ultrium 3 WORM 800-GB data cartridge—C7973W

The HP 50-GB data cartridge in Ultrium 1 format is no longer available. Media from other manufacturers may be used if carrying the official Ultrium logo.

Figure 9.



Conclusions

This paper discusses the new features and the technology of the HP StorageWorks Ultrium 960 tape drive. The Ultrium 3 format delivers capacity, performance, reliability, and industry standards, making it a compelling solution for high-performance data protection solutions. With the additional WORM feature and industry-leading Ultra 320 SCSI interface, HP leads the way in Ultrium technology. The well-established roadmap ensures compatibility across the generations and provides investment protection as well as security of data. The Ultrium logo compliance program ensures an open format with a highly reliable data interchange with all the advantages that it brings. HP again provides a high-technology solution at an affordable price.

References

ECMA Ultrium 316 standard

For more information

<http://www.hp.com/storage>

<http://www.hp.com/go/connect>

<http://www.lto-technology.com>

© 2004 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

5983-0148EN, 12/2004

