Digital Data Storage (DDS) —a Description and Glossary for Customer Engineers

Manual part number C1500-90903

Printed: February 1990
Printed in the UK

Edition 1
E0290

HEWLETT PACKARD

Hewlett-Packard Ltd.
Filton Road, Stoke Gifford, Bristol, BS12 6QZ, UK

Notice

Printing History

Edition 1 . . . . February 1990

Introduction

This document gives a description of DDS-Format drives and the DDS format to provide a background for the CE Service Handbooks for DDS-Format drives. The first part is a continuous description of the format. The second part is in the form of a Glossary, and consists of much the same information arranged differently.

The DDS Format

The Background

Both ½-inch and ¼-inch tape drives use longitudinal recording, in which the data is written in tracks the length of the tape. The track locations are controlled mechanically; this limits the proximity of tracks and hence the areal recording density.

Digital Data Storage (DDS) drives use helical-scan recording, where the recording heads are mounted on a drum which spins at high speed (2000 rpm) on an axis tilted at about 6° to the vertical. The tape wraps 90° around the drum and moves slowly (8 mm/s). The result is that the heads write a series of diagonal tracks across the tape. The heads are aligned on these tracks electronically. The drive compares signals from neighboring tracks and adjusts the drum and tape speeds through a closed-loop servo to ensure accurate tracking. This precision allows areal recording densities one to two orders of magnitude better than longitudinal recording.

Helical scan has been used commercially for video recording in formats such as VHS, BetaMax, and so on. DAT (Digital Audio Tape) took the technology into the audio market, using digital recording and sophisticated error-correction techniques. These features, coupled with the high capacity, low cost, and compact size of the tape cassettes, made DAT an attractive proposition for computer data storage. In 1988, Hewlett-Packard and Sony developed the DDS format, which builds on DAT to produce a tape format suitable for computer data.

DDS-format drives use a similar mechanism to DAT drives. On top of this, the data is organized in groups, rather than the continuous stream used in DAT, and extra error-correction and data protection facilities are provided to ensure the data integrity needed in computer mass storage. DDS also allows for a fast-search facility of up to 200 times the normal read/write speed. This means that with a suitable host and software it is possible for a drive to access any file on a tape within approximately 20 seconds.

The Tape Layout

Cassette tapes can be used by a DDS-format drive as a single partition (a single data space) or as two partitions. The two partitions can be used independently, so the effect is of having two tapes in the drive. The drive can reformat a tape from one partition to two partitions and the other way round, though all data on the tape will be lost in the process.

One-Partition Tape

Figure 1 shows the organization of a tape as a single partition in the DDS format.

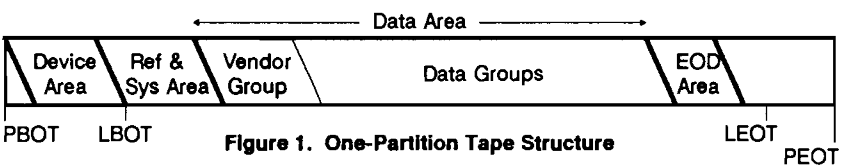

Figure 1. One-Partition Tape Structure

This is the default when a new tape is used.

PBOT and PEOT are the physical beginning and end of tape, where the leader tapes join the magnetic tape. LBOT is the logical beginning of tape, where recorded information begins. LEOT is a point on the tape which warns the drive that PEOT is approaching.

Device Area The Device area is long enough for the tape to be wrapped round the drum, and the drum to accelerate to the reading speed while still in the area. It is during these operations that the tape is most at risk, so the area is blank except for a test region at the end.

The test region is used for read and write tests.

Ref & Sys Area The Reference and System area holds logs of usage and soft errors. It also defines LBOT (Logical Beginning of Tape) which provides a reference point for the tape.

Data Area In the Data area, data is structured in *groups*, described in the section "The Structure of a Group".

Vendor Group The Vendor Group is the first group in the Data area. It contains information about the drive that formatted the tape.

EOD Area The End-of-Data area specifies the point on the tape where the drive stopped writing data. If there was data on the tape before, and there is less new data than there was old, there will be two EODs on the tape – the new one and the old one. The only EOD the drive takes notice of is the one nearer the beginning of the tape, which will be the most recent EOD. All data after that point is ignored by the drive.

Two-Partition Tape

The DDS format allows a tape to be structured as two partitions. An example of use is for Partition 1 (nearer the beginning of the tape) to contain a directory, and Partition 0 (nearer the end of the tape) to contain data. In order to find a particular record in the data, the drive reads the directory in Partition 1 to find where the record is in Partition 0. It can then use the fast-search facility to move to the record and read it. Not all operating systems support this option.

The layout of a two-partition tape is shown in figure 2.

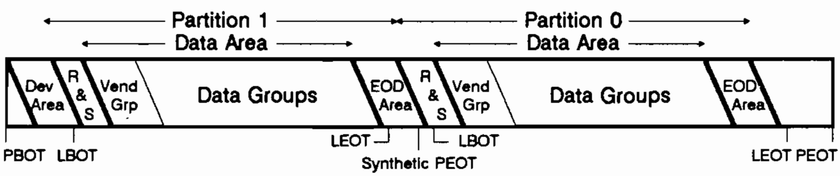

Figure 2. Two-Partition Tape Structure

Each partition has its own Reference and System area, Vendor Group, and EOD area. It follows that each partition has its own logs of tape usage and soft errors. These are needed because the pattern of use of the partitions may be very different.

The Structure of a Group

Data is stored in the Data area in groups with a fixed capacity of 22 *frames* of data (see the next section "The Structure of a Frame"). A group includes a variable-length Index, which is included in the capacity, and an optional extra frame containing Error Correction (ECC) bytes.

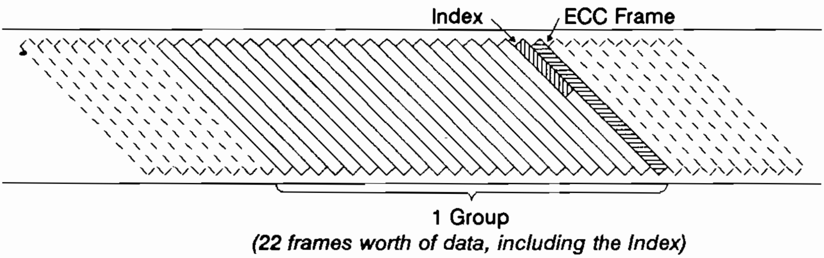

Figure 3. Group Structure

Index The Index stores the position of different records (logical collections of data) within the group. It also holds the position of Filemarks and Save-Set marks. The exact meaning of these marks depends on the host computer and the software it is using, but they provide markers within the Data area and the drive can search or fast-search to them. Save-set marks are hierarchically superior to filemarks. An analogy, or possible use, might be in a list of companies, divided into countries. Each company entry could be a series of records – name, town, telephone number, and so on. A filemark might be used to separate companies. A save-set mark could separate countries.

The Index grows or shrinks as necessary depending on the number of records and marks, so that the maximum amount of space within the group is used for data.

ECC Frame The ECC frame stores Level 3 Error Correction Codes (C3 ECC), and is optional. C3 ECC is described in the Error Correction section later.

A group has a fixed capacity, but not necessarily a fixed length. For example, if a frame is found to be bad when writing data, it is rewritten using read-after-write error correction. In this case there will be 23 frames in the group, but only 22 frames-worth of data.

The Structure of a Frame

The drum on a DDS tape drive has four heads – two write-heads (A and B) and two read-heads (A' and B'). A frame consists of two adjacent tracks, one recorded by the A head on the drum, the other by the B head. The heads are set at different angles, so that the tracks are recorded at different azimuth angles, as shown in figure 4.

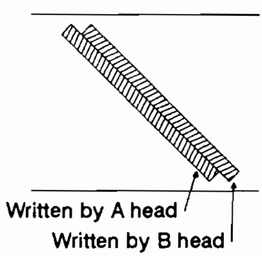

Figure 4. Azimuth Angles

The tape head is wider than a track. This enables tracks to be written overlapping slightly, so that no tape is wasted between tracks. It also enables the head to follow the track exactly. For example, if the A' head is reading the A track, the data on that track produces a strong signal, since the data is at the same azimuth as the head. The head also picks up signals from the B tracks on either side, but these are very weak because the tracks are only partially covered by the head and are recorded at a different azimuth angle. The drive compares the strength of these two weak signals, and alters the drum and tape speeds until they are equal, thus centering the head on the track. The different azimuth angles also minimize cross-talk between neighboring tracks.

The structure of a frame is shown in figure 5. Each track contains 196 blocks, each containing 288 bits. In the User Data area, each block contains 256 bits of data, and the other 32 bits are used for control, address, parity and synchronizing information.

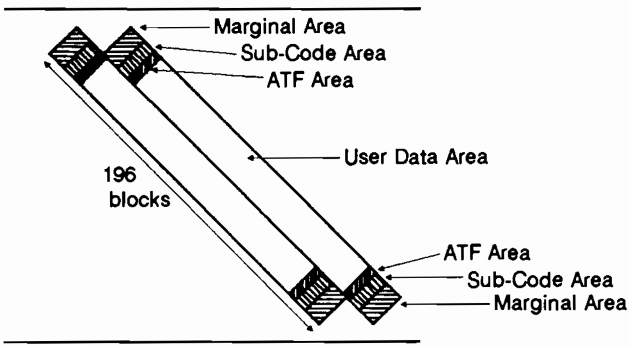

Figure 5. The Structure of a Frame

Marginal Area This carries 11 blocks of a pattern at the data code frequency, which helps synchronize the data stream which follows.

Sub-Code Area The Sub-Code area consists of 8 blocks which carry several copies of the Group Count, Filemark Count, Save-Set Count, and Record Count of the group. The Sub-Code area is read during fast search, and enables the drive to find a particular record or mark. Fast search can be made at 200 times the normal read-write speed, enabling access to any part of the tape typically within 20 seconds.

ATF Area The ATF (Automatic Track Following) area is 5 blocks long and contains signal patterns repeating every four tracks. It is while reading this area that the head picks up patterns from the neighboring tracks, and by balancing the signals, centers itself on the track.

Error-Correction

DAT provides error-correction facilities which enable almost flawless audio reproduction. These facilities are retained in the DDS format, together with extra techniques to bring data integrity up to the standard desired for computer use.

The audio techniques consist of two levels of Error Correction Code, C1 ECC and C2 ECC, which combine with a process called *interleaving* to correct burst errors up to 792 bytes long. The method of writing data using these techniques is as follows:

> Data bytes are subjected to C1 ECC and C2 ECC coding algorithms which produce parity bits and bytes. Two C1 ECC parity bits are attached to each data byte to make a 10-bit symbol. These data symbols are then interleaved. They are rearranged so that consecutive symbols are written far away from each other. The C2 ECC parity bytes—also made into 10-bit symbols with C1 ECC parity bits—are written in the same frame as the data.
>
> When the frame is read, the data is first de-interleaved, so that it is returned to its original order. Any large gap in the data, caused by a burst error from a head clog, perhaps, will now appear as a series of shorter gaps. The parity information is then used to reconstruct these gaps. C1 ECC is used primarily to correct random errors (single bits), while C2 ECC corrects burst errors (sequences of corrupt bits, perhaps extending to several bytes). Without interleaving, C2 ECC can correct up to 6 bytes error. With it, the burst error can be up to 792 bytes long.

DDS adds various correction techniques, including a third level of Error Correction Code (C3 ECC), Randomizing, Read-After-Write, and N-Group writing. The following table lists possible sources of error, and the techniques which can best cope with each category:

Error Source		Correction Technique
Inherent tape defects	Dropouts	C1, C2, RAW, Media specification
	Particles and scratches	C1, C2, RAW, Read retry, Media spec.
	Other surface faults	RAW, Media specification
	Width fluctuation	RAW, Media specification
Damage after data is written	Helical damage	C3, Track checksum
	Longitudinal damage	C1, C2
	Transverse damage	C1, C2
	Other surface damage	Media specification, RAW
Head clogs		RAW, Read retry, C3, Track checksum Head cleaning virtually eliminates head clogs
Mechanical & tape design & production errors	Tracking errors	Drive design
	RF level fluctuations	Media specification
	Jitter and noise	Drive design
	Intersymbol interference	Randomizing
Tape degeneration with age		Media specification

C3 ECC C3 ECC allows any two frames in a group to be corrected. The data is stored in an extra frame attached to the end of each group. It is optional, and particularly valuable if Read-After-Write and N-Group writing are not used.

RAW (Read-after-write) Adding two extra heads to the DAT head drum, so that there are two for reading and two for writing, allows data to be read immediately after it is written. Frames are repeated if they are defective, and so RAW provides an effective means of sparing bad areas on the tape.

N-Group Writing Each group is repeated, so that there are *N* copies on the tape. This is a simple way of increasing data reliability, but it reduces the effective capacity of the tape and performance is sacrificed. *N* is software configurable.

Randomizing Data is recoded so that it provides a consistently high signal level. This is necessary for RAW to work well, where the actual signal level is compared against a threshold value. Randomizing also reduces intersymbol interference. It is estimated to improve the uncorrected raw bit error rate by a factor of 10, with no decrease in performance.

Track Checksum The checksums are recorded in the Sub-Code areas of the tracks. They are useful for detecting *drop-ins*, when previously recorded data remains intact in the middle of new data because of a head clog or other fault. Drops-ins cannot be detected through C1 and C2 ECC. Track checksums can detect which tracks are faulty, and these can be reconstructed through C3 ECC.

Glossary

ATF area

Automatic Track Following is a method of ensuring that the head is in the center of the track being read. The ATF areas are 5 blocks long at the beginning and end of each track on a tape, with the User Data area between them. The head is wider than a track, so while it is reading the ATF area of a track it also reads the ATF areas of the two adjacent tracks. By comparing the content and strength of the signals from the ATF areas, the servo system adjusts the drum and tape speeds so that the read-head passes directly over the center-line of the track being read.

Automatic Track Following

—*see* "ATF area"

Azimuth

The angle at which a recording head is set on the drum, relative to the direction of motion. The two write-heads are set at different azimuth angles (±20°) and write alternate tracks on the tape. The two read-heads have the same azimuth angles as the write-heads. Each read-head reads tracks with its own azimuth. Because the head is wider than the track, the head also picks up signals from the adjacent tracks. These, however, are of the opposite azimuth, so the signals are relatively very weak, and cross-talk is minimal.

Block

Each track consists of 196 blocks. A block is 288 bits, containing 256 bits of data, whether user data, sub-code data, or ATF data. The other 32 bits carry control, address, parity and synchronizing information.

Burst error

A series of contiguous bytes on the tape which are incorrect.

C1 ECC

Level 1 Error Correction Code is a Reed-Solomon code which can correct errors in up to two separate bytes on a track, or burst errors up to four bytes long. The code is stored in two extra bits stored with each byte, making a 10-bit *symbol*.

C2 ECC

Level 2 Error Correction Code is a Reed-Solomon code capable of correcting errors in up to three separate bytes on a track, six bytes erasure error or 792 bytes burst error within a single track. The code is stored on the same track as the data.

C3 ECC

Level 3 Error Correction Code is an optional extra error-correction facility added in the DDS format to allow any two tracks in a group to be corrected, and is only used when a raw data error is too big to be corrected by C1 and C2 ECC. The C3 code is stored in an extra frame at the end of the 22 frames of data in each group.

Checksum

—*see* "Track checksum"

DAT

Digital Audio Tape is a helical scan recording technique for audio recording. It uses 3.81 mm wide tape stored in a cassette measuring 73x54x10.5 mm. A standard 60 m tape cassette can hold up to two hours of digitally recorded music with a quality comparable with that of Compact Discs. The DDS format builds on top of the DAT format to produce a recording technique for computer data.

Data area

The collection of groups on a tape which hold user data. The area includes the Vendor Group.

DDS

Digital Data Storage format is a recording format which builds on the DAT format to support the storage of computer data. It was developed originally by Hewlett-Packard and Sony as an industry standard.

Device area

An area of the tape between PBOT and LBOT (physical and logical beginning of tape). It is long enough so that the tape can be wrapped round the drum and the drum can accelerate up to reading speed while still within the area. It is during these times that the tape is most prone to damage, so the Device area is blank except for a test region at the end.

Drop-In

If a head clog, for example, occurs while writing over data which is already on the tape, the drive may fail both to record the new data and to obliterate the previous data. This patch of unerased data is called a *drop-in*. C1 and C2 ECC cannot detect drop-ins, but track checksums can determine which tracks are faulty and these can be reconstructed through C3 ECC.

Drum

The four recording heads—two read-heads and two write-heads—are mounted on a rotating drum. The drum is inclined at 6% to the vertical and spins at 2000 rpm. The tape is wrapped round the drum in a 90° arc and moves at 8.15 mm/s in the same direction. This gives a relative head to tape writing speed of 3.13 m/s (about 123 ips).

ECC frame

The ECC frame is a frame attached to the end of each group on a tape, and holds Level 3 Error Correction Codes (C3 ECC). The use of C3 ECC is optional within the DDS format.

ECC

Error Correction Code—*see* "C1 ECC", "C2 ECC", and "C3 ECC"

EOD area

The End-of-Data area specifies the point on tape where the drive stopped writing data. If there was data on the tape before the drive started writing, and if the drive has not reached the end of the old data, there will be two EOD areas on the tape. The drive only takes notice of the EOD area nearer to beginning of the tape, which will the most recent EOD. All data after that point is invalid and ignored by the drive. A two-partition tape has a separate EOD area for each partition.

Error correction

The DDS format has ten error correction and detection facilities—C1, C2 and C3 ECC, read-after-write, N-Group writing, error monitoring and reporting (through the use of Error Logs), media specification, retry on read, data randomizer, and track checksums.

Error Correction Code

—see "C1 ECC", "C2 ECC" and "C3 ECC"

Fast-search

Given a suitable host and software, DDS-format drives have the potential for fast-searching a tape for a particular record, filemark or save-set mark. Record, filemark and save-set marks counts are held in the Sub-Code areas on the tape, and during fast-search, the drive reads these at up to 200 times the normal read/write speed.

Filemark

A mark which the host tells the drive to write, and for which the drive can search, sometimes using the fast-search facility. The meaning of the mark is decided by the host software; a filemark does not necessarily separate files. A filemark is hierarchically inferior to a save-set mark.

Filemark Count

The number of filemarks written since the beginning of the partition up to and including the current group. The filemark count is stored in the Index for the group.

Frame

Two adjacent tracks on a tape, one each written by the two write-heads and read by the read-heads with the corresponding azimuth.

Group

A set of frames with a fixed capacity of 22 frames-worth of data. A group contains one Index, and can contain several records, partial records, filemarks, and save-set marks. If C3 ECC is being used, this is stored in an extra frame attached to the end of the group.

Head clog

A head-clog occurs when particles from the tape or outside the drive adhere to the head gap on a read-head or a write-head and obstruct the reading or writing of data. The particles will often become dislodged again with continued use. Head clogs cause burst errors. C3 ECC is the principal error-correction technique to overcome them.

Index

Information stored at the end of each group which describes the contents of the group, for example, specifying where records begin and end, where filemarks occur, the filemark count, and so on. The Index grows and shrinks as necessary, so that there is always the maximum possible space in the group for user data. Every group except the Vendor group contains an Index.

Interleaving

Interleaving is the process of shuffling the order of data bytes before writing them to tape, so that consecutive bytes are recorded far away from each other. This minimizes the impact of any burst error, so that C1 and C2 ECC have the maximum chance of recovering the data.

Layout of tape

—see "Tape layout"

LBOT

Logical Beginning of Tape is the point where the drive starts writing data to the tape and provides a reference point for the tape. LBOT lies between the

Device area and the Reference and System area on the tape. On a two-partition tape, each partition has its own LBOT.

LEOT

Logical End of Tape is a point on the tape some distance before PEOT, which enables the drive to warn the host that it must stop writing data. On a two-partition tape, both partitions have an LEOT.

Marginal area

Marginal areas are 11 blocks long at the ends of each frame. They carry a pattern at data clock frequency which helps synchronize the data stream which follows.

Multiple group writing

—*see* "N-Group writing"

N-Group writing

N-Group writing is an optional technique for avoiding errors, in which each group of data is written several times so that there are *N* consecutive copies on the tape. *N* can take values from 0 through 7 and is software selectable. This is a simple way of improving data integrity, but speed and capacity are sacrificed through the repetition of data.

One-partition tape

A tape with a single Data area. By default, a new tape will be treated by the drive as a one-partition tape. An existing two-partition tape can be reformatted to one partition, but all data will be lost in the process.

Partitions, two

—*see* "Two-partition tape"

PBOT

Physical Beginning of Tape is the point where the leader tape joins the magnetic tape. Following it is the Device area and then LBOT.

PEOT

Physical End of Tape is the point where the magnetic tape joins the trailer tape. On a two-partition tape, Partition 1 has a synthetic PEOT which mimics the function of the real PEOT in Partition 0.

Random error

An error which only involves a single data bit.

Randomizing

Randomizing is the process of recoding data bytes before they are written to tape in order to provide a consistent signal strength. This aids data integrity because the read-after-write process looks for frames with abnormal signal strengths and uses this as a criterion for rewriting the frame.

RAW

—*see* "Read-after-write"

Read-after-write

Read-after-write improves data integrity by reading frames immediately after they are written and rewriting frames if they appear to be in error.

Record

A logical collection of data as defined by the host computer. The DDS format supports variable-length records, so that several short records can be written in one group, and a long record can span several groups.

Retry on read

An error-correction facility in which the drive attempts to re-read data up to seven times.

Reference and System area

The Reference and System area of a tape holds logs of usage and soft errors. It also defines LBOT. In a two-partition tape, each partition has its own Reference and System area.

Save-set mark

A mark written to tape to which the drive can fast-search. The meaning of a save-set mark depends entirely on the host software. Save-set marks are hierarchically superior to filemarks.

Single Partition tape

—*see* "One-partition tape"

Soft error

If the drive can recover data through C3 ECC or read-after-write, the error is considered a soft error. An error from which the data cannot be reconstructed is called a "hard" error.

Structure of Frame

—*see* "Frame"

Structure of Group

—*see* "Group"

Sub-Code area

Areas of eight blocks found at either end of each track between the Marginal areas and the ATF areas. The Sub-Code area holds several copies of the Group Count, Filemark Count, Save-Set Count, and Record Count of the group. During fast-search, the drive reads only these Sub-Code areas and so can quickly move to a specific group, filemark, save-set mark or record.

Symbol

A byte plus two error-correction (C1 ECC) bits is called a symbol, and so data is written to tape in symbols rather than bytes. For simplicity, the word "byte" has often been used in this document where "symbol" would be more accurate.

Synthetic PEOT

—*see* "PEOT"

System area

—*see* "Reference and System area"

Tape heads

DDS-format drives have four tape heads on the rotating drum, two for reading and two for writing. DAT drives only have two which both read and write. By using four in DDS-format drives, read-after-write error correction is made possible, and this is one of the most powerful error-correction methods for computer data.

Tape layout

Tapes may be structured with one or two partitions. In either case, the tape begins with the Device area, which is used for drum spin-up and testing. Each partition begins with LBOT and a Reference and System area, which holds tape logs. After this is a Data area, with a Vendor Group (containing details of the drive that formatted the partition) followed by Data Groups.

Each partition then has an EOD area, defining the end of valid data in the partition, and ends with PEOT or synthetic PEOT. Some distance before PEOT or synthetic PEOT is LEOT which warns the host that PEOT is approaching.

Tape log

The Tape log contains details of the history of the tape, and the total numbers of groups written, of RAW retries, of groups read, of C3 ECC retries, and of loads. The log is copied into RAM when the tape is loaded, updated as the tape is used, and loaded back into the Reference and System area when the tape is unloaded. Note that if the tape is write-protected, the Tape log cannot be updated.

Track

A track is a strip of data written diagonally across the tape by a single write-head as the drum rotates through 90°. For the next 90° rotation, no data is written. During this time the tape has advanced so that for the third 90° of rotation the other write-head on the drum writes the next track of data, overlapping the first by about ⅓ of its width. The result is tracks about 14 μm wide inclined at about 6° to the length of the tape, with no gap between tracks. The two heads are also set at different azimuths so that the azimuth of tracks alternates along the tape. Each track has a User Data area of 148 blocks, flanked on each side by ATF areas of 5 blocks, Sub-Code areas of 8 blocks, and Marginal areas of 11 blocks. A pair of adjacent tracks is called a *frame*.

Track checksum

Checksums for each track are recorded in the Sub-Code areas. They are useful for detecting *drop-ins*, where data previously on the tape is not erased when new data is written, probably due to a head-clog. C1 and C2 ECC cannot detect drop-ins, but track checksums can, and C3 ECC can then reconstruct the missing data.

Two Partition tape

Tapes can be formatted as two partitions, which can be written and read independently. Each partition has its own Reference and System area, Vendor Group and EOD area. The partition nearest to the beginning of the tape is Partition 1, and might be used as a directory for data in Partition 0. The drive could use the directory to find the location of particular data in Partition 0 and then fast-search to the record.

User Data area

Each track has a User Data area of 148 blocks, which include C1 and C2 ECC data.

Vendor Group

The Vendor Group is the first group of the Data area in any partition. It holds details of the drive which formatted the partition or first wrote it, such as the manufacturer and the date.