The FAT File System

As described in the Disk Geometry readings, data is stored on the surface of a platter in ***sectors*** and ***tracks***. Tracks are concentric circles, and sectors are pie-shaped wedges on a track, like this:

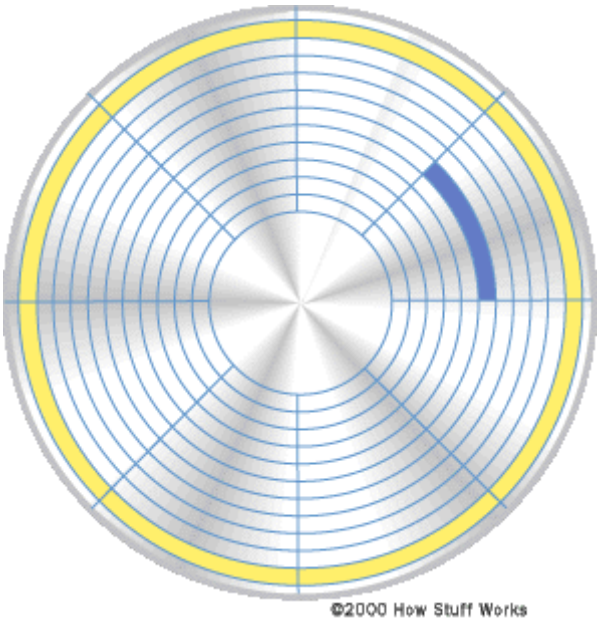

A normal 1.4 Mbyte PC floppy disk has 80 tracks, containing 18 sectors, and may hold 720 Kbytes of data on each side (80 x 18 x 512B) To access data on the disk, the controller hardware needs track, surface and sector values, (referred to as CHS, Cylinder Head & Sector values). But at a higher level, files are located by cluster number. A ***cluster*** is just four contiguous 512 Byte sectors. In FAT-12 systems, each cluster is given a unique 12 bit address. The disk directory, which is held on the first two tracks, starting at sector 19, allows file names to be paired with the number of the initial file cluster.

The floppy disk layout is as follows:

- sector 0 boot track
- sectors 1-9 first File Allocation Table (FAT)
- sectors 10-18 second FAT
- sectors 19-32 disk directory
- sectors 33-2879 data area

On floppy disks, the boot track is one (512 byte) sector long, each FAT is 9 sectors and the directory is 14 sectors. When a disk is formatted, it shows only 698 sectors available. Double-sided floppies also have 1 sector boot tracks.

Directory Structure

The root or main directory is in the 14 consecutive sectors mentioned above. A sub-directory (folder or path) is actually a file on the disk. While the root directory is limited in the maximum number of files it can take, since files need not be contiguous, a subdirectory is not so limited.

Each directory entry is 32 bytes long. There are 16 entries per sector, laid out (the 32 bytes of a directory entry shown in the top row)

0-7	8-10	11	12-21	22-23	24-25	26-27	28-31
File name	exten	attributes	Reserved space	Time	date	Starting cluster	File size

Here is the directory listing:

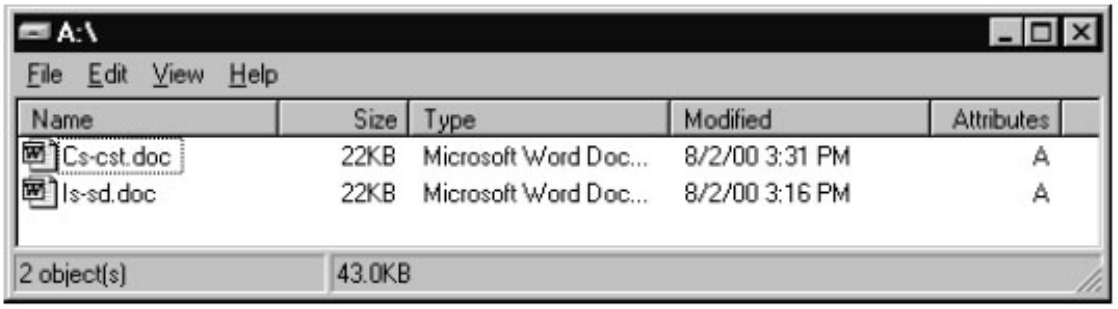

Here are the bytes shown in hex:

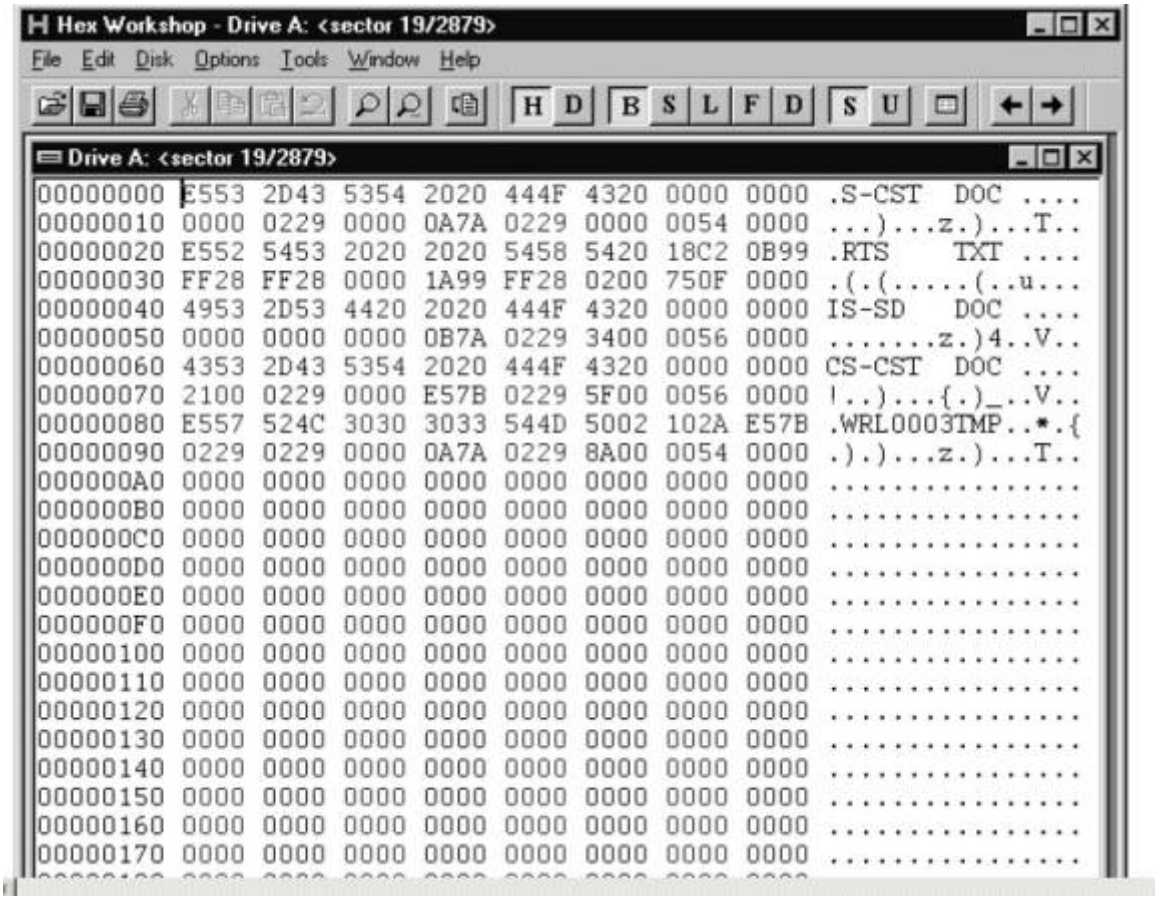

File name - ASCII characters padded with blanks. The first byte of the name field can indicate status:

- $00 entry unused,
- $E5 file erased.
- $2E subdirectory file.

Name Extension - up to 3 ASCII chars, such as TXT, C, DOC, TMP.

Attr ibutes - describe how the files are treated by the system. If the bit is 0, the attribute is off, if 1 it is on or used:

- bit 0 read only
- bit 1 hidden
- bit 2 system file
- bit 3 volume label, not a file
- bit 4 subdirectory

- bit 5 archive bit
- bits 6 and 7 are unused at present

The subdirectory attribute indicates a path folder rather than a file, although they are stored in the same manner. A subdirector y is structured in the same manner as a root directory except that has no fixed size and can grow as large as necessary.

The archive bit is set if a file has been changed since its last backup. (Used normally only with hard disks) Attributes are marked on the disk infor mation screen by a single letter. If an 'x' appears below that letter, then the attribute is set.

Time and date stamps - record the last change of a file.

time stamp:

- bits 0-4 are seconds/2 (0 - 29)
- bits 5-10 are minutes (0 - 59)
- bits 11-15 are hours (0 - 23)

date stamp:

- bits 0-4 are the date (1-31)
- bits 5-8 are the month (1-12)
- bits 9-15 are the year (add 1980 to get the correct value)
-

Starting cluster - an Intel format integer : that is, the low byte comes before the high byte in memory order. It is simply the cluster number where the file begins (there are four 512 byte sectors in each cluster). For where the file goes after that, see the FAT (as described below).

File size - contains the size on disk in bytes but the actual size in memory may be smaller if the last cluster in the FAT is reached before this many bytes are read.

File Allocation Table - FAT

The disk FAT contains many linked chains indicating the clusters associated with each file; the FAT has a one-to-one correspondence with the disk clusters after the first two entries. Each subsequent entry refers to the next cluster on the disk. The FAT and directory size is the same on single and double sided floppies.

On floppies, each 12 bit number in the FAT records a cluster belonging to a particular file. 12 bits means an entry can range from zero to 4096. On a hard disk, each is 16 bits, or a (byte-swapped) integer (which allows a range to 65535).

If a FAT entry is zero, then the associated cluster is unused and available. If 4081 to 4087, then the cluster is unuseable due to a formatting error (usually a bad sector). If the entry contains 4088 to 4095 (FFF), then the cluster contains the end of a file and does not point to another.

In the figure below, the rectangle shows the disk directory, which is always in Sectors 19-32 (see the beginning of this reading). Each row is a directory entry using the 32 bytes described earlier. This figure shows a single file on the disk called "TEST.DAT". It shows that the first cluster is cluster 4, so when a user opens TEST.DAT the disk head goes first to cluster 4 (as shown by the CHS (Cluster Head Sector) pointer leading into the disk in the figure). The disk controller reads the first data in TEST.DAT from cluster 4, then it moves to cluster 7 because the FAT for the disk has a pointer to cluster 7 stored in position 4. This tells the disk head to move from cluster 4 to cluster 7. Similarly when it is done with cluster 7, it follows the pointer and goes to cluster 6 and reads data, finally from cluster 6 it goes to cluster 10 and read data. The pointer in cluster 10 is FFF so it knows the file has been completely read.

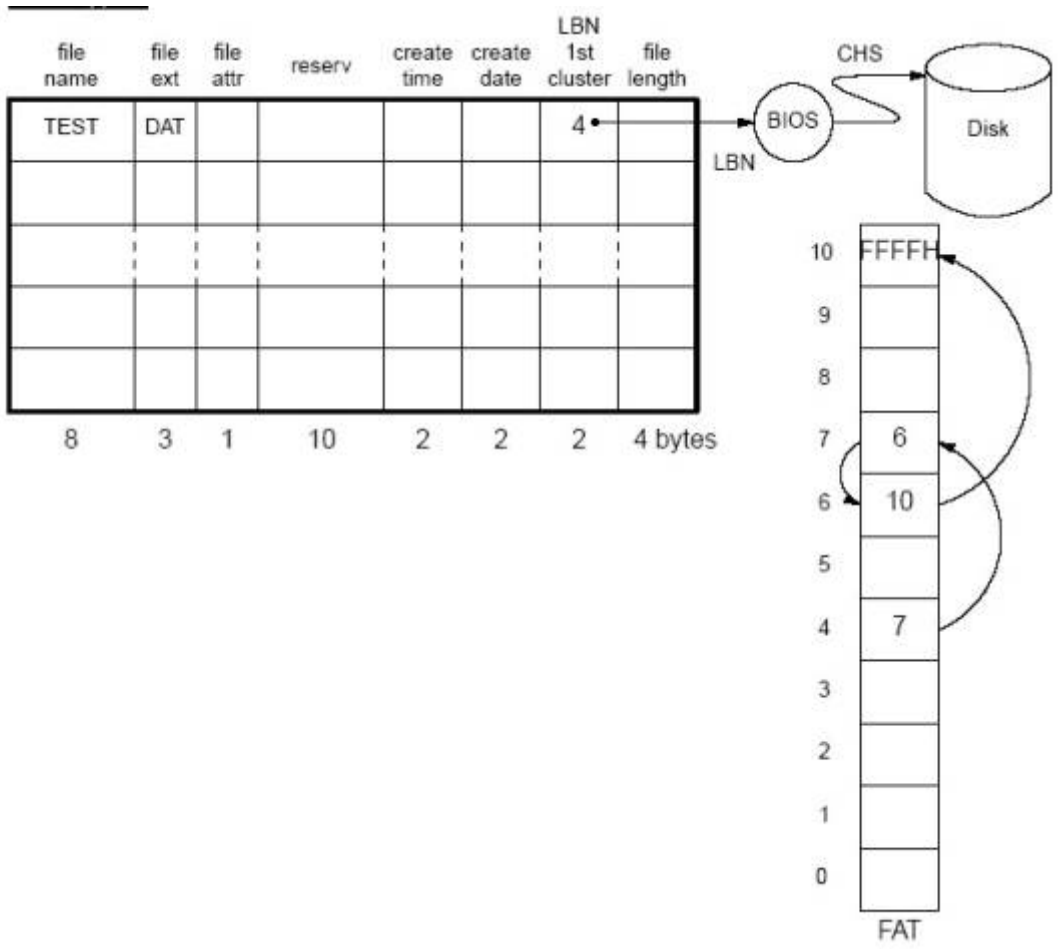

Each file system can define clusters and tables differently (this reading is how FAT does it), but they all basically do the same thing – lay out how files are placed on sectors.