

How It Works: Partition Tables

[\[Home\]](#) [\[Q&A\]](#) [\[History\]](#) [\[Help\]](#) [\[SATA\]](#) [\[SCSI/1394\]](#) [\[HowItWorks\]](#) [\[Search\]](#)

[\[Hale's Test Software\]](#) [\[ATACT\]](#) [\[ATADEMO\]](#) [\[ATAMDT\]](#) [\[ATADRVR\]](#) [\[ADVDRVR\]](#)

[\[HIW Topics\]](#) [\[CHSxlat\]](#) [\[DOSboot\]](#) [\[FAQ\]](#) [\[FNF\]](#) [\[MBR\]](#) [\[OS2boot\]](#) [\[PartTable\]](#)

All About Partition Tables and Sectors and FDISK

PLEASE READ THIS: The following information is only the most basic of the basics. The most important thing to remember about partition tables is that there are no published standards. Hale no longer attempts to keep up with all the silly and stupid things that OS developers are doing in partition tables.

See the following for more information on partition tables:

- [Partition Types by Andries Brouwer](#)

Here is the the basics of how partition tables work, of course, subject to change as some as someone releases a new and popular FDISK program:

- [PARTITION SECTOR/RECORD/TABLE BASICS](#)
- [PARTITION TABLE ENTRY FORMAT](#)
- [PARTITION TABLE RULES](#)
- [PARTITION NESTING](#)
- [PARTITION TABLE LINKING](#)
- [EXAMPLE 1](#)
- [EXAMPLE 2](#)
- [EXAMPLE 3](#)
- [NOTES](#)
- [PARTITION TYPE CODES](#)

PARTITION SECTOR/RECORD/TABLE BASICS

FDISK creates all partition records (sectors). The primary purpose of a partition record is to

hold a partition table. The rules for how FDISK works are unwritten but so far most FDISK programs (DOS, OS/2, WinNT, etc) seem to follow the same basic idea.

First, all partition table records (sectors) have the same format. This includes the partition table record at LBA 0 (cylinder 0, head 0, sector 1) -- what is known as the Master Boot Record (MBR). The last 66 bytes of a partition table record contain a partition table and a 2 byte signature. The first 446 bytes of these sectors usually contain a program but only the program in the MBR is ever executed (so extended partition table records could contain something other than a program in the first 466 bytes). See [Master Boot Record](#).

Second, extended partitions are nested inside one another and extended partition table records form a linked list. I will attempt to show this in a diagram below.

PARTITION TABLE ENTRY FORMAT

Each partition table entry is 16 bytes and contains things like the start and end location of a partition in CHS, the start in LBA, the size in sectors, the partition type and the active flag. Warning: older versions of FDISK may compute incorrect LBA or size values. And note: When your computer boots itself, only the CHS fields of the partition table entries are used (another reason LBA does not solve the >528MB problem). The CHS fields in the partition tables are in L-CHS format, see [How It Works CHS Translation](#).

There is no central clearing house to assign the codes used in the one byte type field, however, there is at least one person at both Microsoft and IBM that attempt to keep track of the type codes. Type codes are used to define most every type of file system that anyone has ever implemented on the x86 PC: 12-bit FAT, 16-bit FAT, HPFS, NTFS, etc. Plus, an extended partition also has a unique type code. Refer to [Partition Type Codes](#) for the complete list of known partition type codes is contained in part 2 of this document.

The 16 bytes of a partition table entry are used as follows:

```
+--- Bit 7 is the active partition flag, bits 6-0 are
|
| zero (when not zero this byte is also the drive
| number of the drive to boot so the active
| partition is always found on drive 80H,
| the first hard disk).
|
|
| +--- Starting CHS in INT 13 call format.
|
|
|     +--- Partition type byte.
|
|
|         +--- Ending CHS in INT 13 call format.
```

								+-- Starting LBA.							
								+-- Size in sectors.							
v	---	+	----	v	---	+	----	v				v			
0.	1.	2.	3.	4.	5.	6.	7.	8.	9.	A.	B.	C.	D.	E.	F.
DL	DH	CL	CH	TB	DH	CL	CH	LBA	SIZE				
80	01	01	00	06	0e	be	94	3e000000		0c610900		1st	entry		
00	00	81	95	05	0e	fe	7d	4a610900		724e0300		2nd	entry		
00	00	00	00	00	00	00	00	00000000		00000000		3rd	entry		
00	00	00	00	00	00	00	00	00000000		00000000		4th	entry		

Bytes 0-3 are used by the small program in the Master Boot Record to read the first sector of an active partition into memory. The DH, DL, CH and CL above show which x86 register is loaded when the MBR program calls INT 13H AH=02H to read the active partition's boot sector. See [Master Boot Record](#).

These entries define the following partitions:

1. The first partition, a primary partition DOS FAT, starts at CHS 0H,1H,1H (LBA 3EH) and ends at CHS 294H,EH,3EH with a size of 9610CH sectors.
2. The second partition, an extended partition, starts at CHS 295H,0H,1H (LBA 9614AH) and ends at CHS 37DH,EH,3EH with a size of 34E72H sectors.
3. The third and fourth table entries are unused.

PARTITION TABLE RULES

Keep in mind that there are NO written rules and NO industry standards on how FDISK should work but here are some basic rules that seem to be followed by most versions of FDISK:

1. In the MBR there can be 0-4 primary partitions, OR, 0-3 primary partitions and 0-1 extended partition entry.
2. In an extended partition there can be 0-1 secondary partition entries and 0-1 extended partition entries.
3. Only 1 primary partition in the MBR can be marked active at any given time.
4. In most versions of FDISK, the first sector of a partition will be aligned such that it is at head 0, sector 1 of a cylinder. This means that there may be unused sectors on the track (s) prior to the first sector of a partition and that there may be unused sectors following a partition table sector.

For example, most new versions of FDISK start the first partition (primary or extended) at cylinder 0, head 1, sector 1. This leaves the sectors at cylinder 0, head 0, sectors 2... n as unused sectors. This same layout may be seen on the first track of an extended partition. See [example 2](#) below.

Also note that software drivers like Ontrack's Disk Manager depend on these unused sectors because these drivers will hide their code there (in cylinder 0, head 0, sectors 2... n). This is also a good place for boot sector virus programs to hang out.

5. The partition table entries (slots) can be used in any order. Some versions of FDISK fill the table from the bottom up and some versions of FDISK fill the table from the top down. Deleting a partition can leave an unused entry (slot) in the middle of a table.
6. And then there is the hack that some newer OS's (OS/2 and Linux) use in order to place a partition spanning or passed cylinder 1024 on a system that does not have a CHS translating BIOS. These systems create a partition table entry with the partition's starting and ending CHS information set to all FFH. The starting and ending LBA information is used to describe the location of the partition. The LBA can be converted back to a CHS, most likely a CHS with more than 1024 cylinders. Since such a CHS can not be used by the system BIOS, these partitions can not be booted or accessed until the OS's kernel and hard disk device drivers are loaded. It is not known if the systems using this hack follow the same rules for the creation of these type of partitions.

There are NO written rules as to how an OS scans the partition table entries so each OS can have a different method. For DOS, this means that different versions could assign different drive letters to the same FAT file system partitions.

PARTITION NESTING

What do I mean when I say the partitions are nested within each other? Lets look at this example:

M = Master Boot Record (and any unused sectors on the same track)

E = Extended partition record (and any unused sectors on the same track) pri = a primary partition (first sector is a boot sector) sec = a secondary partition (first sector is a boot sector)

```
|-----the entire disk-----|
|
|M[pri]                          |
|      E[sec]----rest of 1st ext part-----|
|              E[sec]----rest of 2nd ext part-----|
```

The first extended partition is described in the MBR and it occupies the entire disk following the primary partition. The second extended partition is described in the first extended partition record and it occupies the entire disk following the first secondary partition.

PARTITION TABLE LINKING

What do I mean when I say the partition records (tables) form a linked list? This means that the MBR has an entry that describes (points to) the first extended partition, the first extended partition table has an entry that describes (points to) the second extended partition table, and so on. There is, in theory, no limited to out long this linked list is. When you ask FDISK to show the DOS logical drives it scans the linked list looking for all of the DOS FAT type partitions that may exist. Remember that in an extended partition table, only two entries of the four can be used (rule 2 above).

And one more thing... Within a partition, the layout of the file system data varies greatly. However, the first sector of a partition is expected to be a boot sector. A DOS FAT file system has: a boot sector, first FAT sectors, second FAT sectors, root directory sectors and finally the file data area. See [OS2 Boot Sector](#).

EXAMPLE 1

A disk containing four DOS FAT partitions (C, D, E and F):

```
|-----the entire disk-----|
|
|M[---C:---]
|           E[---D:---]--rest of 1st ext part-----|
|                   E[---E:---]--rest of 2nd ext part--|
|                                   E[-----F:-----]|
```

EXAMPLE 2

So here is an example of a disk with two primary partitions, one DOS FAT and one OS/2 HPFS, plus an extended partition with another DOS FAT:

```
|-----the entire disk-----|
|
```

```

| M[pri 1 - DOS FAT] |
| [pri 2 - OS/2 HPFS] |
| E[sec - DOS FAT] |

```

Or in more detail (n is the highest cylinder, head or sector number allowed in the indicated field of the CHS)...

```

+-----+
CHS=0,0,1 | Master Boot Record containing |
| partition table search program and |
| a partition table |
| +-----+ |
| | DOS FAT partition description | | points to
CHS=0,1,1 | +-----+ | points to CHS=a
| | OS/2 HPFS partition description | |
| +-----+ |
| | unused table entry | |
| +-----+ |
| | extended partition entry | | points to CHS=b
| +-----+ |
+-----+
CHS=0,0,2 | the rest of "track 0" -- this is | :
to | where the software drivers such as | : normally
CHS=0,0,n | Ontrack's Disk Manager or Micro | : unused
| House's EZ Drive are located. | :
+-----+
CHS=0,1,1 | Boot sector for the DOS FAT | :
| partition | | : a DOS FAT
+-----+ | : file
CHS=0,1,2 | rest of the DOS FAT partition | : system
to | (FAT table, root directory and | :
CHS=x-1,n,n | user data area) | :
+-----+
CHS=x,0,1 | Boot sector for the OS/2 HPFS | :
| file system partition | | : an OS/2
+-----+ | : HPFS file
CHS=x,0,2 | rest of the OS/2 HPFS file system | : system
to | partition | :
CHS=y-1,n,n | | :
+-----+

```

CHS=y,0,1	Partition record for the extended partition containing a partition record program (never executed) and a partition table	
	+-----+ DOS FAT partition description	points to CHS=b+1
	+-----+ unused table entry	
	+-----+ unused table entry	
	+-----+ unused table entry	
	+-----+	
CHS=y,0,2 to CHS=y,0,n	the rest of the first track of the extended partition	: normally : unused :
	+-----+	
CHS=y,1,1	Boot sector for the DOS FAT partition	: : a DOS FAT : file
	+-----+	
CHS=y,1,2 to CHS=n,n,n	rest of the DOS FAT partition (FAT table, root directory and user data area)	: system : :
	+-----+	

EXAMPLE 3

Here is a partition record from an extended partition (the first sector of an extended partition). Note that it contains no program code. It contains only the partition table and the signature data.

```

OFFSET 0 1 2 3 4 5 6 7 8 9 A B C D E F *0123456789ABCDEF*
000000 00000000 00000000 00000000 00000000 *.....*
000010 TO 0001af SAME AS ABOVE
0001b0 00000000 00000000 00000000 00000001 *.....*
0001c0 8195060e fe7d3e00 0000344e 03000000 *.....}....4N....*
0001d0 00000000 00000000 00000000 00000000 *.....*
0001e0 00000000 00000000 00000000 00000000 *.....*
0001f0 00000000 00000000 00000000 000055aa *.....U.*

```

PARTITION TYPE CODES

Hale no longer maintains a list of partition type codes. The following is probably the best list kept by anyone:

- [List of partition identifiers for PCs](#)

This page was last updated on 06 May 2002.