by
Jeff Prosise

Tutor

THE DOS PARTITIONING SCHEME

Can you shed some light on the MS-DOS partitioning scheme and how it works? What's the difference between a primary DOS partition and an extended DOS partition? And why can't you install more than one primary partition on a hard disk?

Bob Mattes
Edwards, California

■ **THE DOS PARTITIONING SCHEME: Here's how MS-DOS divides your hard disk.**

Partitioning is a means of dividing a single hard disk into several logical devices. DOS treats separate partitions on a hard disk as if they were separate hard disks. Having a 60MB hard disk with two 30MB partitions is the same as having two 30MB hard disks. You see them as drives C: and D:; DOS takes care of the details associated with accessing them as such.

Partitioning was originally developed so that different operating systems could share the same hard disk. The MS-DOS partitioning scheme permits a hard disk to be divided into as many as four partitions. Prior to Version 3.3, only one of these partitions could be used by DOS. Version 3.3 eliminated this restriction with the introduction of the extended DOS partition, which is a collection of logical volumes grouped together in a linked list format. Embedded within each volume is the address of the next volume in the chain. These volumes are treated just like normal disk partitions, even to the extent that they can be formatted with any DOS or OS/2-compatible file system.

The key to the MS-DOS partitioning scheme lies in the partition table FDISK sets up in the first physical sector of a hard disk. The partition table is a 64-byte structure that defines where the partitions are located on a hard disk, which one is bootable, and what file system each one was formatted with.

Figure 1 shows the format of the partition table, stored at offset 1BEh in the first sector on head 00h, cylinder 00h. It contains four 16-byte entries—one for each of four possible partitions. The first byte in each entry indicates whether or not the corresponding partition is bootable. If it's bootable, it should contain the bootstrap code needed to get the PC up and running. For bootable partitions, FDISK writes the value 80h to this slot; for others, it writes 00h. Only primary partitions can be designated as bootable. When FDISK marks a DOS partition as such, it scans the rest of the partition table and zeros any boot indicators it finds that aren't already zeroed, ensuring that DOS is the operating system that receives control at start-up.

The bytes at offsets 01h through 03h and 05h through 07h define the physical limits of the partition. The first three designate the starting head, sector, and cylinder numbers, and the second three the ending head, sector, and cylinder numbers. The sector number is only a 6-bit value. The extra bits aren't wasted, though; the upper 2 bits in the sector field are joined with the byte in the cylinder field to form a 10-bit cylinder number.

The byte at offset 04h, known as the system indicator byte, specifies what type

THE MS-DOS PARTITION TABLE FORMAT

	00h	01h	02h	03h	04h	05h	06h	07h	08h-0Bh	0Ch-0Fh
01BEh	B	H	S	C	I	H	S	C	RelSec	NumSec
01CEh	B	H	S	C	I	H	S	C	RelSec	NumSec
01DEh	B	H	S	C	I	H	S	C	RelSec	NumSec
01EEh	B	H	S	C	I	H	S	C	RelSec	NumSec

Field definitions

Offset	Symbol	Meaning
00h	B	Boot indicator (1 byte) 00h = not bootable 80h = bootable
01h	H	Starting head number (1 byte)
02h	S	Starting sector number (6 bits)
03h	C	Starting cylinder number (10 bits*)
04h	I	System indicator (1 byte) 00h = unused 01h = 12-bit FAT 04h = 16-bit FAT 05h = extended partition 06h = 32-bit FAT 07h = OS/2 HPFS
05h	H	Ending head number (1 byte)
06h	S	Ending sector number (6 bits)
07h	C	Ending cylinder number (10 bits*)
08h	RelSec	Number of preceding sectors (4 bytes)
0Ch	NumSec	Length of partition in sectors (4 bytes)

* Formed by adding the upper 2 bits of the previous byte to the 8 bits of the current byte. The extra bits become bits 8 and 9.

Figure 1: The partition table, located in the first sector of hard disks prepared with DOS or OS/2, defines the physical location of up to four partitions, what file system each was formatted for, and which partition is bootable.

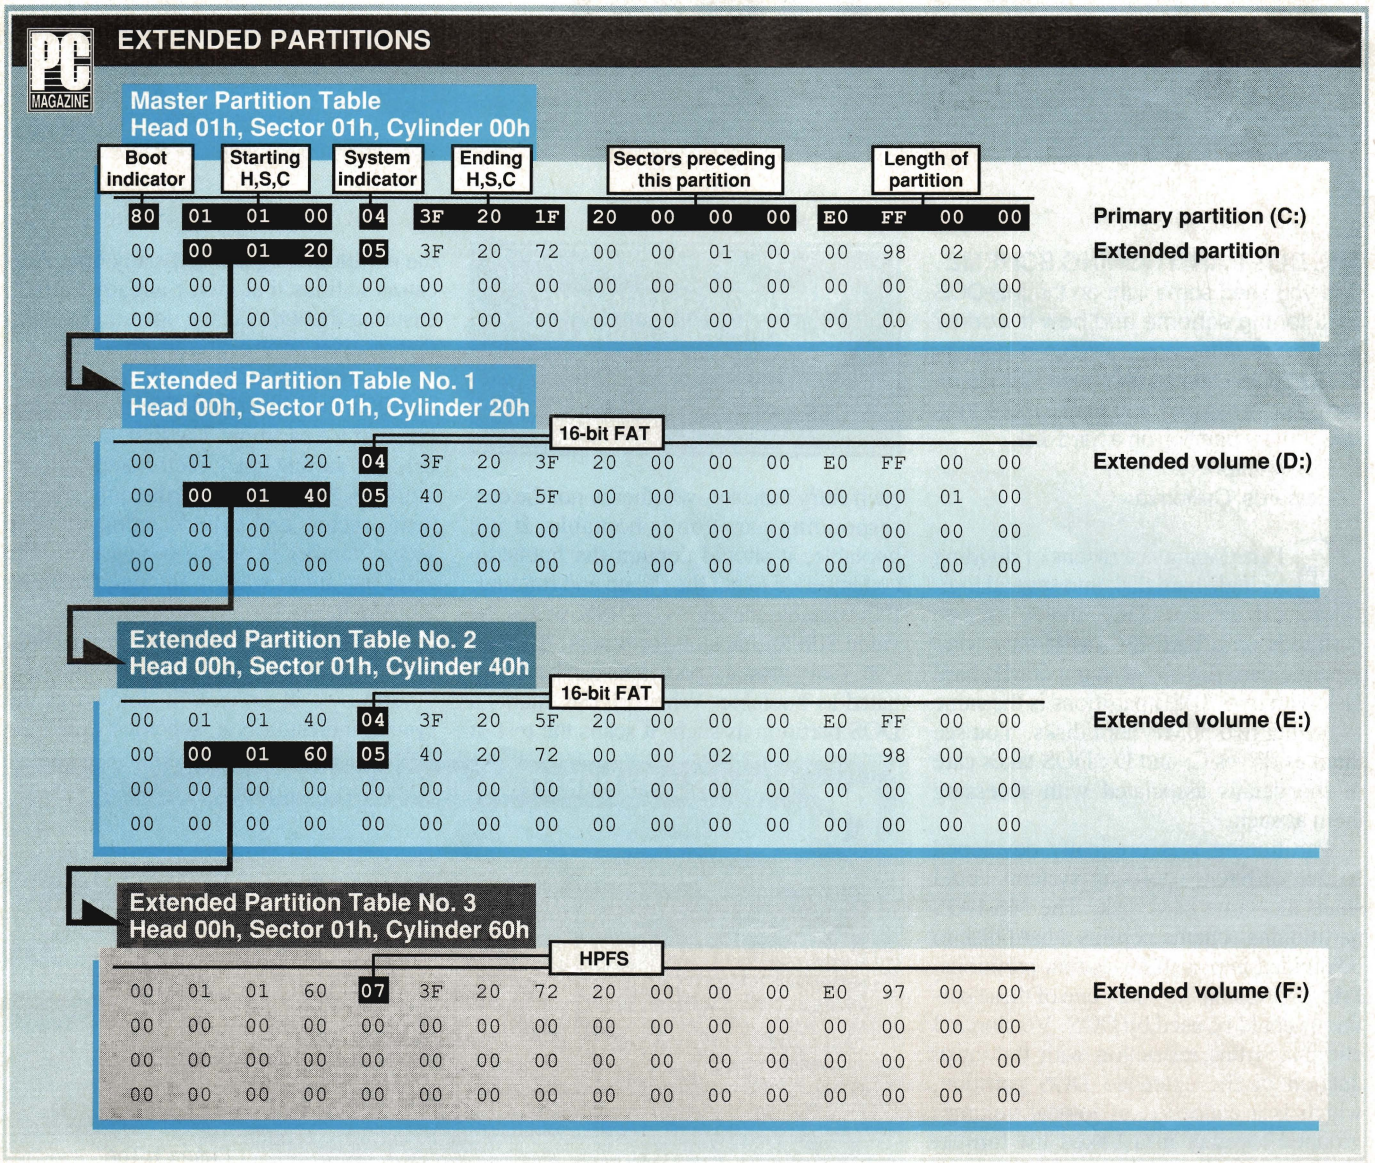

Figure 2: The chain of partition table entries for a 120MB hard disk divided into one primary partition and one extended partition. The extended partition contains three logical volumes. The final volume is formatted for OS/2's High Performance File System.

of file system the partition was formatted for. A value of 00h means the partition is unused. A 01h identifies a DOS partition that uses 12-bit FAT (file allocation table) entries, while a 04h identifies a DOS partition that uses 16-bit FAT entries. An extended DOS partition is marked with the value 05h. A 06h indicates that the corresponding partition is a huge (greater than 32MB) partition set up by DOS 4.0 or OS/2. And 07h is the value assigned to partitions formatted with OS/2's HPFS (High Performance File System).

The key to the MS-DOS partition scheme is the 64-byte structure set up by FDISK.

What happens if you install an operating system that mistakenly uses one of these reserved values to identify its own partition? Fortunately, this doesn't happen often, because most vendors of PC-based operating systems are aware of DOS's conventions and steer clear. But when it does, havoc can result. At least one version of Xenix I've seen writes an ID value of 06h to the partition table. If you're running DOS 3.*x* or earlier, you're okay—DOS will recognize that it's somebody else's partition. But DOS 4.0 and OS/2 will think it's a huge partition with 32-bit FAT entries. Attempting to access it could result in damaged data. The last thing you want to do is to run CHKDSK on it.

The last two fields in the partition table entry, at offsets 08h and 0Ch, are double-word entries that hold the starting sector

number relative to the first sector of the disk and the total length of the partition in sectors, respectively. Although these values are somewhat redundant (DOS could easily calculate them based on other values in the partition table) having them stored in this form is helpful because it eliminates some of the work involved in accessing the partition.

During start-up, DOS scans the partition table for partitions formatted with DOS-compatible file systems and assigns logical drive letters on a first-come, first-served basis. In most versions of DOS, partition table entries are read from front to back, so the entry at offset 1BEh becomes drive C:, the entry at 1CEh becomes drive D:, and so on. Some DOS resellers, however, tweak the code to read entries from back to front. In this case, the partition whose entry appears at offset 1EEh becomes drive C:, the one at 1DEh becomes drive D:, and so on. If there is more than one hard disk, DOS assigns drive letters to primary DOS partitions before allocating room for extended partitions.

The entry for a primary DOS partition is a straightforward one. The partition corresponds to one logical volume. DOS assigns a drive letter to it at start-up. Its characteristics are spelled out in black and white in the partition table, and locating it on-disk is a simple matter of consulting the partition table to find out what head, sector, and cylinder it starts on.

Utilizing extended partitions enables DOS 3.3 users to make more-efficient use of their hard disks by allowing them to bypass the four-partition limit found in the MS-DOS primary partitioning scheme.

PARTABLE.SCR — COMPLETE LISTING

```
A
MOV  AX,0201          ;Function 02H
MOV  CX,0001          ;Cylinder 00H, sector 01H
MOV  DX,0080          ;Head 00H, drive 80H
MOV  BX,0200          ;ES:BX -> CS:0200
PUSH CS
POP  ES
INT  13               ;Read the sector
RET                   ;Exit

G
D 03BE 03FD
Q
```

Figure 3: This short DEBUG script reads the first sector of a hard disk. The instruction D03BE displays the master partition table. Be sure to type in the code exactly as it is shown; a miscue could render your hard disk unreadable.

Figure 2 shows the chain of partition table entries for a 120MB hard disk formatted with a primary DOS partition and an extended DOS partition. Actually, the names are somewhat misleading: the primary "DOS" partition in this case is used exclusively for OS/2. A better way to describe these partitions would be as a primary 16-bit FAT partition and an extended partition.

The disk's master partition table contains two entries. The first defines the primary DOS partition, which starts at head 01h, sector 01h, cylinder 00h and ends at head 3Fh, sector 20h, cylinder 1Fh—a total of 65,504 sectors, or 32.8MB. The extended partition spans the range from head 00h, sector 01h, cylinder 20h, to head 3Fh, sector 20h, cylinder 72h, consuming almost 85MB worth of storage space.

Finding the first extended volume in the extended partition requires that we read the starting sector spelled out for the extended partition in the master partition table and decipher the extended partition table stored there. The first extended volume is a type 04h that is 65,504 sectors in length. The second entry describes another extended partition, which contains a second extended volume identical to the first. Proceeding one level deeper yields the final extended volume, a 19MB type 07h formatted for OS/2's HPFS.

I hope the picture is becoming clearer. Extended partitions neatly circumvent the limit of four partitions imposed by the size of the master partition table. They do this by chaining together as many logical volumes as a disk will hold. FDISK limits you to one primary DOS partition, because extended partitions are the preferred mechanism for installing multiple DOS partitions. Had DOS 3.3 users been restricted to primary partitions, they would also have been limited to 120MB per hard disk—far too little to enable DOS to make efficient use of large hard disks.

If you're curious to see what the master partition table on your hard drive looks like, you can use the short DEBUG script shown in Figure 3 to view it. The script actually has you assemble a short machine code program that uses the BIOS Read Sector service to read the hard disk's first physical sector. Once the contents of the sector have been read into memory, the instruction D 03BE displays the partition table. Be very careful to type in the script exactly as it is shown, especially the line MOV AX,0201; a miscue could render the hard disk unreadable.

ASK THE TUTOR

The Tutor solves practical problems and explains techniques for using your hardware and software more productively. Questions about DOS and systems in general are answered. To have your questions answered here, write to Tutor, *PC Magazine,* One Park Avenue, New York, NY 10016, or upload them to PC MagNet (see the "By Modem" sidebar in the Utilities column). We're sorry, but we cannot answer questions individually. ■