# Macintosh: File System Specifications & Terms (10/96)

Article Created: 13 August 1991
Article Reviewed/Updated: 31 October 1996

TOPIC -----------------------------------------------------------

What are the limitations of the Hierarchical File System (HFS) on the Macintosh? Are there practical limits after which performance suffers? This article explains some of the terminology involved when speaking about volumes and the file system. This article also explains why different HFS structures on different hard drives can cause small files to take up more space on a larger drive than they do on a smaller drive.

DISCUSSION -----------------------------------------------------

The Hierarchical File System Specifications
==========

Before discussing hard drive block allocation, and the intricacies of determining allocation block sizes and such, let’s first review the Hierarchical File System, better known as HFS.

Below are some specifications and associated terminology used to describe the Macintosh HFS structure. The following numbers apply to both System 6 and System 7, and are current as of this writing.

Volume
---
A volume is either a full disk, or a section of a disk, partitioned into separate parts. If you partition a single drive then each partition is considered a volume.

• The maximum volume size under System 6 and System 7 is 2 gigabytes. System 7.5 increased that limit to 4 GB and System 7.5.2 (and later) increases that limit to 2 terabytes on some computers, including:
  - PowerBook 190, 1400, 5300 series and Duo 2300c/100.
  - Macintosh Performa 5400 and 6400 series and 6360/160
  - Power Macintosh 5400, 6400, 7200, 7500, 7600, 8500, and 9500 series
  - Workgroup Server 7250/120 and Workgroup Server 8550 series

• The maximum number of volumes is limited only by the amount of available system memory.

• The maximum file size is 2 gigabytes.

• The maximum number of files on a volume is 65,536.

• The maximum number of files in a folder is 32,767.

• The maximum size of the data fork in a file is 2 gigabytes.

• The maximum size of the resource fork in a file is 16 megabytes.

Logical block
---
A logical block is a unit of drive space composed of up to 512 bytes. A logical block is numbered from 0 to n, n being the last block on the volume – not necessarily the hard disk. Take the volume size, divide it by 512 bytes, and you have the number of logical blocks.

Allocation block
---
An allocation block is a unit of storage on a volume, composed of one or more logical blocks. The larger the volume, the more logical blocks comprise one allocation block. The maximum number of allocation blocks per volume is 65,536; most volumes have slightly less.

In both the Macintosh and DOS environment, the maximum number of blocks on a driver is 65,536 because both Operating Systems address the allocation blocks with a 16-bit address. Drives larger than 512 MB cannot use a block size of 8K or less because there just aren’t enough addresses. Thus, if a 2 GB drive is one Macintosh partition, the smallest file size allowed is 32K. (If you save a TeachText file with one character in it, it would take up 32K of disk space.) This means that the size of your hard drive determines the minimum size of each file.

A non-empty file fork always occupies at least one allocation block, no matter how many bytes of data the file fork contains. If the data fork and resource fork of a file contain data, the file uses two allocation blocks. If all the files are only one allocation block long, then there can be as many files on the volume as there are allocation blocks for the volume. However, when a file is longer than one allocation block, the total number of possible files decreases. For example, on a volume with 65,535 allocations block, you can have:

```
   65,536 files, each 1 allocation block long.
   32,768 files, each 2 allocation blocks long.
     2048 files, each 32 allocation blocks long.
     1024 files, each 64 allocation blocks long.
```

All of the numbers above refer to the System's record structures. Other elements place tighter restrictions on the actual number of usable files in the root directory. The Standard File dialog uses the List Manager. The List Manager has a limit of 32K of data, which is somewhere between 800 and 900 files. Though more files can be handled in a directory, the List Manager will not display the first 32K.

A second element concerns the performance of the Finder when approaching 800 to 900 files per directory. Even simple operations, such as moving the icon of a file, tend to slow down considerably when this number of files exists in a directory. This limitation changes depending on the speed of the individual Macintosh model.

## Determining Allocation Block Size

Follow these steps to calculate the allocation block size:

1) Take the size of the drive in megabytes and multiply it by 2000 (there are 2000 disk blocks in 1 MB of disk space).

2) Divide that number by 65,536, since the Macintosh cannot have more than 65,536 allocation blocks.

3) Round this number up and multiply it by 512.

The result is the number of bytes in the allocation block. Here is an example using a 230 MB drive:

```
230 x 2000 / 65536 = 7.019
7.019 rounded up = 8
8 x 512 = 4096 bytes
```

So what does this mean to you? The larger the hard drive, the larger the allocation block size, and the more space that is wasted on small files. If you have a large drive with a lot of small files, the hard drive space is being used less efficiently than if most of your files average 32K in size.

Important Note: Since drives come from different manufacturers, it is possible to get different values. This is because a 500 MB drive is not exactly 500 MB but is usually somewhat larger. This could result in different logical block sizes than those listed in the chart (see Allocation Block Size Table below), but the difference should not be more than +/- 512 bytes. This difference is most noticeable on the larger drives. For example, the 500 MB drive actually calculates to 8192 bytes but the 500 MB drive in some Macintosh computers is actually 540 MB, which has a logical block size of 8704, as listed in the chart.

## Allocation Block Size Table

Here is a complete list of allocation block sizes for the various volume sizes, based on the formula described above:

Begin_Table

| Volume Size | Allocation Block Size | Minimum file size |
|---|---|---|

| | | | |
|---|---|---|---|
| 0-32MB | 1 | logical block | 0.5K |
| 33-64MB | 2 | logical blocks | 1.0K |
| 65-96MB | 3 | logical blocks | 1.5K |
| 97-128MB | 4 | logical blocks | 2.0K |
| 129-160MB | 5 | logical blocks | 2.5K |
| 161-192MB | 6 | logical blocks | 3.0K |
| 193-224MB | 7 | logical blocks | 3.5K |
| 225-256MB | 8 | logical blocks | 4.0K |
| 257-288MB | 9 | logical blocks | 4.5K |
| 289-320MB | 10 | logical blocks | 5.0K |
| 321-352MB | 11 | logical blocks | 5.5K |
| 353-384MB | 12 | logical blocks | 6.0K |
| 385-416MB | 13 | logical blocks | 6.5K |
| 417-448MB | 14 | logical blocks | 7.0K |
| 449-480MB | 15 | logical blocks | 7.5K |
| 481-512MB | 16 | logical blocks | 8.0K |
| 513-544MB | 17 | logical blocks | 8.5K |
| 545-576MB | 18 | logical blocks | 9.0K |
| 577-608MB | 19 | logical blocks | 9.5K |
| 609-640MB | 20 | logical blocks | 10.0K |
| 641-672MB | 21 | logical blocks | 10.5K |
| 673-704MB | 22 | logical blocks | 11.0K |
| 705-736MB | 23 | logical blocks | 11.5K |
| 737-768MB | 24 | logical blocks | 12.0K |
| 769-800MB | 25 | logical blocks | 12.5K |
| 801-832MB | 26 | logical blocks | 13.0K |
| 833-864MB | 27 | logical blocks | 13.5K |
| 865-896MB | 28 | logical blocks | 14.0K |
| 897-928MB | 29 | logical blocks | 14.5K |
| 929-960MB | 30 | logical blocks | 15.0K |
| 961-992MB | 31 | logical blocks | 15.5K |
| 993-1024MB | 32 | logical blocks | 16.0K |
| 1025-1056MB | 33 | logical blocks | 16.5K |
| 1057-1088MB | 34 | logical blocks | 17.0K |
| 1089-1120MB | 35 | logical blocks | 17.5K |
| 1121-1152MB | 36 | logical blocks | 18.0K |
| 1153-1184MB | 37 | logical blocks | 18.5K |
| 1185-1216MB | 38 | logical blocks | 19.0K |
| 1217-1248MB | 39 | logical blocks | 19.5K |
| 1249-1280MB | 40 | logical blocks | 20.0K |
| 1281-1312MB | 41 | logical blocks | 20.5K |
| 1313-1344MB | 42 | logical blocks | 21.0K |
| 1345-1376MB | 43 | logical blocks | 21.5K |
| 1377-1408MB | 44 | logical blocks | 22.0K |
| 1409-1440MB | 45 | logical blocks | 22.5K |
| 1441-1472MB | 46 | logical blocks | 23.0K |
| 1473-1504MB | 47 | logical blocks | 23.5K |
| 1505-1536MB | 48 | logical blocks | 24.0K |
| 1537-1568MB | 49 | logical blocks | 24.5K |
| 1569-1600MB | 50 | logical blocks | 25.0K |
| 1601-1632MB | 51 | logical blocks | 25.5K |
| 1633-1664MB | 52 | logical blocks | 26.0K |
| 1665-1696MB | 53 | logical blocks | 26.5K |

```
1697-1728MB      54 logical blocks         27.0K
1729-1760MB      55 logical blocks         27.5K
1761-1792MB      56 logical blocks         28.0K
1793-1824MB      57 logical blocks         28.5K
1825-1856MB      58 logical blocks         29.0K
1857-1888MB      59 logical blocks         29.5K
1889-1920MB      60 logical blocks         30.0K
1921-1952MB      61 logical blocks         30.5K
1953-1984MB      62 logical blocks         31.0K
1985-2016MB      63 logical blocks         31.5K
2017-2048MB      64 logical blocks         32.0K
2049-2080MB      65 logical blocks         32.5K
2081-2112MB      66 logical blocks         33.0K
2113-2144MB      67 logical blocks         33.5K
2145-2176MB      68 logical blocks         34.0K
2177-2208MB      69 logical blocks         34.5K
2209-2240MB      70 logical blocks         35.0K
2241-2272MB      71 logical blocks         35.5K
2273-2304MB      72 logical blocks         36.0K
2305-2336MB      73 logical blocks         36.5K
2337-2368MB      74 logical blocks         37.0K
2369-2400MB      75 logical blocks         37.5K
2401-2432MB      76 logical blocks         38.0K
2433-2464MB      77 logical blocks         38.5K
2465-2496MB      78 logical blocks         39.0K
2497-2528MB      79 logical blocks         39.5K
2529-2560MB      80 logical blocks         40.0K
2561-2592MB      81 logical blocks         40.5K
2593-2624MB      82 logical blocks         41.0K
2625-2656MB      83 logical blocks         41.5K
2657-2688MB      84 logical blocks         42.0K
2689-2720MB      85 logical blocks         42.5K
2721-2752MB      86 logical blocks         43.0K
2753-2784MB      87 logical blocks         43.5K
2785-2816MB      88 logical blocks         44.0K
2817-2848MB      89 logical blocks         44.5K
2849-2880MB      90 logical blocks         45.0K
2881-2912MB      91 logical blocks         45.5K
2913-2944MB      92 logical blocks         46.0K
2945-2976MB      93 logical blocks         46.5K
2977-3008MB      94 logical blocks         47.0K
3009-3040MB      95 logical blocks         47.5K
3041-3072MB      96 logical blocks         48.0K
3073-3104MB      97 logical blocks         48.5K
3105-3136MB      98 logical blocks         49.0K
3137-3168MB      99 logical blocks         49.5K
3169-3200MB     100 logical blocks         50.0K
3201-3232MB     101 logical blocks         50.5K
3233-3264MB     102 logical blocks         51.0K
3265-3296MB     103 logical blocks         51.5K
3297-3328MB     104 logical blocks         52.0K
3329-3360MB     105 logical blocks         52.5K
3361-3392MB     106 logical blocks         53.0K
```

| | | |
|---|---|---|
| 3393-3424MB | 107 logical blocks | 53.5K |
| 3425-3456MB | 108 logical blocks | 54.0K |
| 3457-3488MB | 109 logical blocks | 54.5K |
| 3489-3520MB | 110 logical blocks | 55.0K |
| 3521-3552MB | 111 logical blocks | 55.5K |
| 3553-3584MB | 112 logical blocks | 56.0K |
| 3585-3616MB | 113 logical blocks | 56.5K |
| 3617-3648MB | 114 logical blocks | 57.0K |
| 3649-3680MB | 115 logical blocks | 57.5K |
| 3681-3712MB | 116 logical blocks | 58.0K |
| 3713-3744MB | 117 logical blocks | 58.5K |
| 3745-3776MB | 118 logical blocks | 59.0K |
| 3777-3808MB | 119 logical blocks | 59.5K |
| 3809-3840MB | 120 logical blocks | 60.0K |
| 3841-3872MB | 121 logical blocks | 60.5K |
| 3873-3904MB | 122 logical blocks | 61.0K |
| 3905-3936MB | 123 logical blocks | 61.5K |
| 3937-3968MB | 124 logical blocks | 62.0K |
| 3969-4000MB | 125 logical blocks | 62.5K |
| 4001-4032MB | 126 logical blocks | 63.0K |
| 4033-4064MB | 127 logical blocks | 63.5K |
| 4065-4096MB | 128 logical blocks | 64.0K |

End_Table

## Volume organization of HFS

The first two logical blocks (labeled 0 and 1) of a volume are the boot blocks. This is where the information for mounting the volume is stored.

The third logical block (labeled 2) is the Master Directory Block, or MDB for short. This block contains part of the data structure of a flat directory volume. It contains the volume information and the volume allocation block map. This block is where the information for the hard disk, such as number of files in the directory and the last time the drive was initialized, is stored.

Logical blocks 3 to x (see the following table) contain the volume bitmap. This block is a data structure containing a sequence of bits, one bit for each allocation block. The volume bitmap stores a reference to every piece of data that is in the allocation blocks and indicates whether the block is allocated or free for use. Volume bitmaps exist both on hierarchical directory volumes and in memory.

Volumes may have as few as 32,768 allocation blocks and as many as 65,536 allocation blocks. This table defines what logical block the volume bitmap ends on based on the number of allocation blocks:

Begin_Table

| Allocation Blocks | Last Volume Bitmap Logical Block |
|---|---|
| > 61,440 | 18 |

| | |
|---|---|
| > 57,344 | 17 |
| > 53,248 | 16 |
| > 49,152 | 15 |
| > 45,056 | 14 |
| > 40,960 | 13 |
| > 36,864 | 12 |
| > 32,768 | 11 |

End_Table

Allocation blocks begin after the volume bitmap. Contained in the allocation blocks are the catalog, extents and data files.

The catalog file is a list of all files and folders stored in a volume. The catalog file maintains the relationships between the files and directories on a hierarchical directory volume. It corresponds to the file directory on a flat directory volume. The catalog file is organized and accessed using a B- tree structure. The files are arranged in alphabetical order evenly balanced on the tree so that finding a "Z" doesn't take any longer than finding an "A". This structure is the glue that keeps the catalog file together.

The extents tree file contains the locations of all the files on a volume (An extent is a series of contiguous allocation blocks). The extents tree file is where the information (such as where to find file, and how many extents a file is divided into) about the data files you have created is stored. Any given file you create may be broken up into multiple extents. When extents are linked together, behind the scenes, with information from the extents tree file, the appearance is of one data file.

The next, and largest, section of the volume contains all the actual data files and applications which are referenced using the above files.

The next to last block on the hard disk contains the alternate master directory. This alternate is a backup to the MDB kept on logical block 2. It is used when the file manager determines that the MDB is corrupt and needs to be rewritten with the correct information.

The very last block is empty. It is used to check for bad sections of the hard disk.

## What You Can Do

So, can you do anything to decrease the file size on a large hard drive? One solution is to partition larger drives into smaller partitions, or logical drives, each with a maximum of 65,536 blocks.

Apple's Drive Setup can be used to partition qualified Apple hard drives and some removable devices as listed in the Drive Setup Guide file. Unsupported hard drives will show up in the device list by name, however, if you select an unsupported hard drive you will get the message, "Cannot modify a disk in an unsupported drive."

Drive Setup is supported on all of the Power Macintosh computers. This includes the original Power Macintosh computers and the new PCI-based Power Macintosh computers. Drive Setup is also supported on 68LC040 processor-based systems that have IDE drives installed. Drive Setup is not supported with the Power Macintosh Upgrade Card. Drive Setup 1.0.2 is available from online services.

There are third parties that offer a solution for creating multiple Macintosh partitions on Macintosh systems not supported by Drive Setup. A few are:

- Hard Disk Toolkit (HDT) by FWB Software, Inc.
- SilverLining by LaCie Ltd.
- Micronet Utility by MicroNet Technology
- MicroTech Utility by MicroTech International
- Drive 7 by Casa Blanca Works, Inc.

This article provides information about a non-Apple product. Apple Computer, Inc. is not responsible for its content. Please contact the vendor for additional information.

The Tech Info Library article titled "Locating Vendor Information" can help you search for a particular vendor's address and phone number.

These articles can help you locate the Drive Setup software update mentioned here:

- "Where To Find Apple Software Updates" - Lists online services for free Apple software updates.

- "Obtaining Apple Product Support in the USA" - Lists 800 numbers and online services for software updates, Apple support information, and a subset of the Apple Tech Info Library.

Article Change History:
31 Oct 1996 - Added list of computers supporting more than 4 GB.
08 Sep 1995 - Reviewed for technical accuracy. Added System 7.5.x limits.
24 May 1995 - Reviewed for technical accuracy.



Keywords: sys7,sys75, kalley,review

==========

This information is from the Apple Technical Information Library.

19961101 12:37:21.00

Tech Info Library Article Number: 8647