

File Systems

From OSDev Wiki

Filesystems are the machine's way of ordering your data on readable and/or writable media. They provide a logical way to access the stuff that you have down on disk so that you can read or modify it. Which file system you use depends upon what you want to do with it. For example, Windows uses the Fat32 or NTFS filesystem. If your disk is really huge, then there's no point using Fat32 because the FAT system was designed in the days when nobody had disks as big as we do now. At the same time, there's no point using a NTFS filesystem on a tiny disk, because it was designed to work with large volumes of data - the overhead would be pointless for, say, reading a 1.44m floppy disk.

For details on specific filesystems, browse this list of filesystems.

Contents

- 1 File System Theory
 - 1.1 Indexing Methods
 - 1.1.1 Inodes
 - 1.1.2 FAT
- 2 Workings of File Systems
 - 2.1 Allocation Table
 - 2.2 Separate file and system areas
 - 2.3 Other methods
- 3 Network File Systems
- 4 File systems for OSDevers
 - 4.1 "Beginners" filesystems
 - 4.2 Rolling your own
 - 4.2.1 Guidelines if you do decide to roll your own
 - 4.3 Expert filesystems
- 5 See Also
 - 5.1 Wiki Pages

File System Theory



This page is a work in progress and may thus be incomplete.

Its content may be changed in the near future.

Filesystems

Virtual FileSystems

VFS

Disk filesystems

FAT 12/16/32, VFAT

Ext 2/3/4

LEAN

HPFS

NTFS

HFS

HFS+

MFS

ReiserFS

FFS (Amiga)

FFS (BSD)/UFS

BeFS

BFS

XFS

SFS

ZFS

CD/DVD filesystems

ISO 9660

Joliet

UDF

Network filesystems

NFS

RFS

AFS

Flash filesystems

JFFS2

YAFFS

A filesystem provides a generalized structure over persistent storage, allowing the low-level structure of the devices (e.g., disk, tape, flash memory storage) to be abstracted away. Generally speaking, the goal of a filesystem is allow logical groups of data to be organized into *files*, which can be manipulated as a unit. In order to do this, the filesystem must provide some sort of index of the locations of files in the actual secondary storage. The fundamental operations of any filesystem are:

- Tracking the available storage space
- Tracking which block or blocks of data belong to which files
- Creating new files
- Reading data from existing files into memory
- Updating the data in the files
- Deleting existing files

(Perceptive readers will note that the last four operations - Create, Read, Update, and Delete, or CRUD - are also applicable to many other data structures, and are fundamental to databases as well as filesystems.)

Additionally, there are other features which go along with a practical filesystem:

- Assigning human-readable names to files, and renaming files after creation
- Allowing files to be divided among non-contiguous blocks in storage, and tracking the parts of files even when they are *fragmented* across the medium
- Providing some form of hierarchical structure, allowing the files to be divided into *directories* or *folders*
- Buffering reading and writing to reduce the number of actual operation on the physical medium
- Caching frequently accessed files or parts of files to speed up access
- Allowing files to be marked as 'read-only' to prevent unintentional corruption of critical data
- Providing a mechanism for preventing unauthorized access to a user's files

Additional features may be found on some filesystems as well, such as automatic encryption, or journalling of read/write activity.

Indexing Methods

There are several methods of indexing the contents of files, with the most commonly used being *i-nodes* and *File Allocation Tables*.

Inodes

Inodes (information nodes) are a crucial design element in most Unix filesystems: Each file is made of data blocks (the sectors that contains your raw data bits), index blocks (containing pointers to data blocks so that you know which sector is the nth in the sequence), and one inode block.

The inode is the root of the index blocks, and can also be the sole index block if the file is small enough. Moreover, as unix filesystems support hard links (the same file may appear several times in the directory tree), inodes are a natural place to store metadata such as file size, owner, creation/access/modification times, locks, etc.

FAT

The File Allocation Table (FAT) is the primary indexing mechanism for MS-DOS and its descendants. There are several variants on FAT, but the general design is to have a table (actually a pair of tables, one serving as a backup for the first in case it is corrupted) which holds a list of blocks of a given size, which map to the whole capacity of the disk.

Workings of File Systems

There are several common approaches to storing disk information. However, in comparison to memory management, there are some key differences in managing disk media:

- Data can only be written in fixed size chunks.
- Access times are different for different locations on the disk. Seeking is usually a costly operation.
- Data throughput is very small compared to RAM
- Data commonly has to be maintained

Hence some file systems have specialized structures, algorithms, or combinations thereof to improve speed ratings.

Allocation Table

The allocation table is comparable to the bitmap approach. However instead of just having a field free or occupied, it may contain other information. Advantage is that the use of this structure is simple, Disadvantage is that this approach is relatively slow, and a separate set of data is needed to define the which sections are used, and in which order. FAT for example, combines a linked list and an allocation table in the same structure.

Separate file and system areas

Some filesystems keep metadata and actual contents in separate areas on the disk. This makes specific sorts of data easy to find, as well as allowing for a special 'free' area. Downside is that you have to keep track of the boundary or boundaries between these areas, as the usage can differ and the disk has to adapt for that (big files: more data, less metadata; small files: less data, more metadata). This method is used prominently in SFS.

Other methods



This page is a work in progress and may thus be incomplete. Its content may be changed in the near future.

Network File Systems

All these file systems are a way to create a large, distributed storage system from a collection of "back end" systems. That means you cannot (for instance) format a disk in 'NFS' but you instead mount a 'virtual' NFS partition that will reflect what's on another machine. Note that a new generation of File Systems is under heavy research, basing on latest P2P, cryptography and error correction techniques (such as the Ocean Store Project or Archival Intermemory.)

For details on various network file systems look [here](#)

File systems for OSDevers

While you could pick <insert favorite filesystem here> as your OS main filesystem, you might want to consider all options. Commonly, you'll want to have a filesystem that is operational very quickly so that you can concentrate on the rest before implementing a 'real filesystem'

"Beginners" filesystems

There are only four filesystems that are both relatively easy to implement and worth to consider. There is no general recommendation as the choice depends largely on style and OS design. Instead you can read the comparison and make your own educated decision.

FAT

- + Can be read and written by virtually all OSes
- + The 'standard' for floppies
- + Relatively easy to implement
- - Patented by Microsoft. If you wish to use long file names you should pay them.
- - Large overhead
- - No support for large (>4GB) files
- - No support for unix permissions

Ext2

- + Supports large files (with an extension)
- + Supports unix permissions
- + Can be put on floppies
- + Can be read and written from linux
- - Can not natively be read and written from windows (but drivers (<http://www.fs-driver.org/>) are available)
- - Very large overhead
- - Of these four, this is the most complex filesystem

SFS

- + Supports large files
- + By far the easiest to implement
- + Can be put on floppies and harddisks
- + Minimal overhead
- - New, and therefore unsupported. The only operational utility is available for Windows.
- - No support for fragmentation
- - No support for unix permissions

ISO 9660

The defined standard for CDs. If you boot from CD then this is the way to go. If not, don't make it your first filesystem.

Rolling your own

There are many different kinds of filesystems around, from the well-known to the more obscure ones. The most unfortunate thing about filesystems is that every hobbyist OS programmer thinks that the filesystem they design is the ultimate technology, when in reality it's usually just a copy of FAT with a change here and there, perhaps because it is one of the easiest to implement. The world doesn't need another FAT-like filesystem. Investigate all the possibilities before you decide to roll your own.

Guidelines if you do decide to roll your own

- Consider carefully what it will be used for.
- Use a program to figure out the layout. I recommend a spreadsheet. The basic areas needed are:
 - Bootsector. Unless you are booting with UEFI, this is a must. Even then, I'd recommend including it in the specs for compatibility with older file systems. This section should contain at a minimum the disk size, location of the file table, hidden sectors for multiple partition disks, and a version number. I'd leaving plenty of reserved space for features you don't think of. Don't forget to leave space for a jmp instruction and the boot code!.
 - File table. Don't think of this as just a simple table containing a list of files and their locations. One idea I've had and never implemented is, instead of storing files, the system would store file parts, and the file table would list the parts in each file. This would be useful for saving space if many files on the disk are the same or similar (for example, license agreements).
 - Data area. Files will be stored here.
- Create a program to read and write disk images with your filesystem. Parts of this will be portable into the fs driver.
- Implement the fs into your OS.

Expert filesystems

Once you have a beginner's file system under your belt you might want support for more advanced ones. Here are some:

- NTFS - (Windows) New Technologies File System. It's hard to find documentation. Try Apple NTFS (open source) (<http://www.opensource.apple.com/source/ntfs/>) .

See Also

Wiki Pages

I use a Custom Filesystem - What Bootloader Solution is right for me?

Retrieved from "http://wiki.osdev.org/index.php?title=File_Systems&oldid=17309"

Categories: In Progress | OS theory | Filesystems

-
- This page was last modified on 4 December 2014, at 23:02.
 - This page has been accessed 74,465 times.