

For my project I wanted to look at file systems to have a better understanding of what goes on when you try to read or write a file. I also wanted to take a look at some of the advanced functions of file systems, and the different file systems that were popular over time.

What does a file system do?

A file system without any advanced features is very simple: it stores data as files to a storage device, and provides a method to access that data. In addition, most file systems attached metadata to the file, such as the date of creation, the size of the file, and even some conditional metadata such as pixel size of an image.

History of file systems

According to Wikipedia, there are 86 different file systems. I only looked at some of the more popular ones, and the ones that deal with a hard drive or a disk drive.

RT-11

The first file system I looked at was RT-11, a file system written for the RT-11 operating system. It was created in 1973 by DEC. It was a flat file system, it did not allow subdirectories. The filename format was 6.3, where filenames were 6 characters long with a 3 character extension. In this file system, the files were stored in contiguous

blocks, which meant that the file system had to be defragmented frequently. The file size was stored as the number of blocks, not bytes.

CP/M

One of the first file systems was the file system for the operating system CP/M. It was created in 1974 by Gary Kindall. It was the first file system to use filenames in the 8.3 format used by DOS, where the filename can contain up to 8 characters with an extension of 3 characters. This was a non-hierarchical file system, it did not allow subdirectories. However, it supported “user areas” which was an attempt to allow multiple users. With the “user” command, it was possible to change the user, which really just changed the files that appeared in the file system. In addition, the file size was stored as the number of 128-byte records, not the number of bytes.

File Allocation Table (FAT 12)

FAT 12 was developed by Microsoft in 1977 for the Disk Basic operating system. This operating system was the predecessor to the widely used FAT 16 and FAT 32 file systems. The file system was designed for use on floppy disks, and is still used on floppy disks today.

Macintosh File System (MFS)

MFS was developed by Apple in 1984 for the original Mac OS (then simply referred to as System Software). It was a flat file system, however the Finder supported

having folders one level deep (no nested subdirectories). These weren't stored in folders, and in open/save dialog boxes displayed all files without any folders. The filename can contain 255 characters, although the Finder only supported 63 characters, and more recent versions only supported 31.

Hierarchical File System (HFS)

HFS was developed by Apple in 1985 for Mac OS. It was a successor to MFS, which natively supported subdirectories. The file system is also referred to as Mac OS Standard. In HFS, files are referred to by their unique ID instead of filename.

File Allocation Table (FAT 16)

FAT 16 was developed by Microsoft in 1987 as an improvement on FAT 12. It was first included in MS-DOS 3.31, and was originally called *DOS 3.31 Large File System*. FAT 16 was created to increase the file size and partition size limit to allow use on larger hard drives.

High Performance File System (HPFS)

HPFS was developed by IBM and Microsoft in 1988 for OS/2. The file system was designed to improve upon FAT, as is a precursor to NTFS. HPFS used an architecture that kept related files close to each other on the disk to speed up read times. It also stored the root directory in the middle of the drive for faster random

access, as opposed to storing it at the beginning of the drive which most file systems do.

ext

ext was created by Rémy Card in 1992 for Linux. The name is short for Extended File System. It was the first file system created specifically for Linux, and was superseded by ext2.

New Technology File System (NTFS)

NTFS was created by Microsoft in 1993 for Windows NT 3.1. There have been several updates to the file system since then. NTFS has many features, including journalling (described later), native file compression, transparent data encryption, and disk quotas. Volume Shadow Copy Service keeps old versions of files by copying old data to a “shadow copy” when new data is written. NTFS supports mount points like Linux uses, where a separate partition can be mounted as a folder instead of another drive. It also supports symbolic links (soft links) (described later).

ext2

ext2 was created by Rémy Card in 1993 as a successor to ext. The biggest difference between it and ext is that it stores new files in a contiguous block of free space to help prevent fragmentation.

File Allocation Table (FAT 32)

FAT 32 was created by Microsoft in 1996 as an improvement over FAT 16. It increases the maximum file and volume size over FAT 16. FAT 32 was first included in Windows 95 OSR2, and supports long file names (up to 255 characters) under Windows 95 and higher using Virtual FAT.

Hierarchical File System Plus (HFS Plus)

HFS Plus was created by Apple in 1998 as a successor to HFS. HFS Plus is also referred to as Mac OS Extended. HFS Plus increased the maximum file and volume size over HFS. In addition, HFS Plus supports journalling.

ext3

ext3 was created by Stephen Tweedie in 1999 as a successor to ext2. It supports journalling and online growth. It does not provide a defragmentation tool, and cannot be defragmented while online.

ReiserFS

ReiserFS was created by Namesys (specifically the company's founder Hans Reiser) in 2001 for Linux. It was the first file system that supported journalling to be included in the Linux kernel. It was the default file system for many distributions of

Linux, and was the default for SUSE Linux Enterprise until Novell chose to use ext3 in September 2006. ReiserFS became a controversial file system to use after its creator was arrested for murder in September 2006 after his wife went missing, coincidentally when Novell dropped ReiserFS as its default file system. Reiser has since pleaded guilty to the charge and led police to his wife.

ZFS

ZFS was created by Sun Microsystems in 2004 for Solaris and FreeBSD. It originally stood for Zettabyte file system. ZFS is built to use multiple disks to contain a single file system (using storage pools). ZFS has a major focus on data integrity, and doesn't support the use of RAID because it prefers to manage all hard drives.

Additional features of file systems

Journalling

Journalling is a simple way to prevent data. It stores data changes in a journal and applies all changes at once, to prevent the possibility of some changes being written while other changes not. The journal can include all data to be written, or just the metadata.

Symbolic Links and Hard Links

Symbolic links are shortcuts to files in another location. They are followed as if they are the actual files themselves, just as clicking a Windows shortcut or Mac alias. Hard links are links that actually point to the data on the disk. Some file systems support having multiple hard links point to the same data. In file systems which support it, deleting a hard link will only delete the data if there are no other hard links pointing to the data.