

# THE COMPUTER CORNER

## No. 300: LINUX: THE FILE SYSTEM

de Stan Kaplan, WB9RQR, 715 N. Dries Street, Saukville, WI 53080-1664  
[wb9rqr@gmail.com](mailto:wb9rqr@gmail.com)

Before computing even existed, the term *file system* was used to describe storing and retrieving paper documents. After the early 1960's, it began to also be applied to computerized filing. The term then grew to describe a series of rules that dictates how a computer operating system (OS) will store and retrieve data. The FAT file system was used in the DOS and early Windows OS, and later NTFS (New Technology File System) took over for the more recent releases of Windows, including Windows 10 and 11.

Linux supports a number of file systems. However, the most current Linux Mint Cinnamon Version 21 (nickname "Vanessa") uses the **ext4** file system, as do other versions of Linux since 2008. Even Google now uses ext 4 on its Android OS.

Ext4 can handle single files up to 16 tebibytes in size (1 tebibyte =  $1024^4$  bytes in the binary system or 1 terabyte =  $1000^4$  bytes in the metric system). It can also deal with volumes (directories) up to 1 exbibyte ( $1024^6$  bytes binary,  $1000^6$  bytes metric). So we have a file system that can handle huge files and huge volumes of files. The block mapping (extants) is also new, which improves large file performance and reduces fragmentation.

Ext4 is backward-compatible with earlier Linux ext3 and ext2 file systems. There is now no limitation on the number of subdirectories except for size limitations as noted earlier, so you can create up to about 10-12 million subdirectories in a directory, should you be so moved! There are other improvements, as well. Faster file-system checking, fancier ways of allocating data to buffers and groups of blocks, all improve handling of data and reducing fragmentation on disk. It seems likely that defragmenter software will be needed less and less; already some experts say don't even bother since file fragmentation is so low. Another new development is based on the fact that second-based timestamps are no longer sufficient for mission-critical applications, so ext4 provides timestamps measured in *nanoseconds*. They also added a couple of bits to the seconds field of the timestamps; this delays a problem destined to rear its ugly head in the year 2038 for an additional 408 years(!!!).

One improvement that helps reduce fragmentation of files involves delaying the writing of data to disks until it is time to flush the data. That delay increases the risk of data loss, should a system power loss or crash occur before the data has been completely written to disk. To avoid this, the latest versions of ext4 handle such cases the way it was done in ext3, with data flushes more often and data loss risk therefore reduced.

A big plus is that ext4 enjoys full support by other operating systems. Since 2016, Windows had the ability to access ext4, and a commercial product (Paragon's Linux File Systems for Windows) is available for earlier versions of Windows and Windows Server. MacOS can read and write ext2, ext3 and ext4 through extFS for Mac, also by Paragon.

There is no doubt that ext4 has improved features, but some experts say it is not a major advance, that it uses old technology, and is a stop-gap system. Some say to look for Btrfs ("B-tree File System" or "Better File System") for future development. Btrfs is a "copy on write" system in which a clone operation automatically produces a copy-on-write snapshot of a file. On the other hand, ext4 is a "journaling file system", in which the system keeps a journal of where files are located on disk, and other changes that are made to the disk. The latter system helps keep fragmentation of files to a minimum, while the former helps keep data loss to a minimum. Perhaps our grandchildren will be the ones to see which approach wins out! Happy Computing!