

Virtual File System on Linux

- Sudhanshoo Maroo(00D05001)
- Virender Kashyap (00D05011)

Virtual File System on Linux.

What is it ?

VFS is a kernel software layer that handles all system calls related to file systems. Its main strength is providing a **common interface** to several kinds of file systems.

What's Linux VFS's key idea?

For each read, write or other function called, the kernel substitutes the actual function that supports a native Linux file system, for example the NTFS.

File systems supported by Linux VFS

- disk based file systems like ext3, VFAT
- network file systems
- other special file systems like /proc

VFS File Model

Superblock object

- Stores information concerning a mounted file system.
- Holds things like device, blocksize, dirty flags, list of dirty inodes etc.
- Super operations -> like read/write/delete/clear inode etc.
- Gives pointer to the *root inode* of this FS
- Superblock manipulators: mount/umount

File object

- Stores information about the interaction between an open file and a process.
- File pointer points to the current position in the file from which the next operation will take place.

VFS File Model

inode object

- stores general information about a specific file.
- Linux keeps a cache of active and recently used inodes.
- All inodes within a file system are accessed by file-name.
- Linux's VFS layer maintains a cache of currently active and recently used names, called *dcache*

dcache

- structured in memory as a tree.
- each entry or node in tree (*dentry*) points to an inode.
- it is not a complete copy of a file tree

Note : If any node of the file tree is in the cache then every ancestor of that node is also in the cache.

VFS Files Path Lookup

How to reach the file or directory?

- Having in hand the inode of the initial directory, the code examines the entry matching the first name to get the corresponding inode.
- Then the directory file having that node is read from disk and the entry matching the second name is examined to derive the corresponding inode.
- This procedure is repeated for each name included in the path.

The dentry cache considerably speeds up the procedure

File system operations are mostly done at the dcache level , so they are all under kernel **lock**.

References

- <http://www.cse.unsw.edu.au/~neilb/oss/linux-commentary/vfs-1.html>
- http://www.itcentrs.lv/linux/docs/Linux_Kernel_Internals/Linux-Kernel-Internals-3.html
- Bram :
http://www.coda.cs.cmu.edu/doc/talks/linux_vfs/