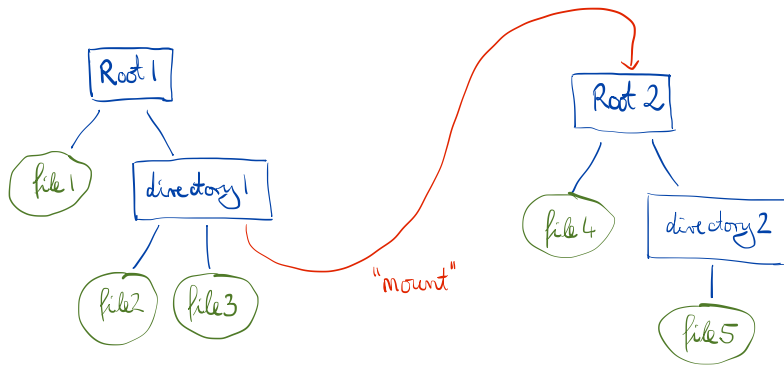


# File systems

Recall: File systems are organized in a tree structure where the files are the leaves and the non-empty directories are the internal nodes.

Mounting: The root of one file system can be "mounted" as a subdirectory of another.



Purpose: present a single tree structure to userspace from physically or technically separate file systems

→ flexibility

→ ease of use once set up properly

## Network file systems

- presented to the user, e.g. by mounting, just as a local file system, but accessed via a network (possibly via the internet)
- For LAN ("Local Area Networks") there are two main options:

NFS (comes from Sun Unix): V3 - simple, fast, does not use a secure protocol (still in use if clients can be trusted)

V4 - complex, more like CIFS

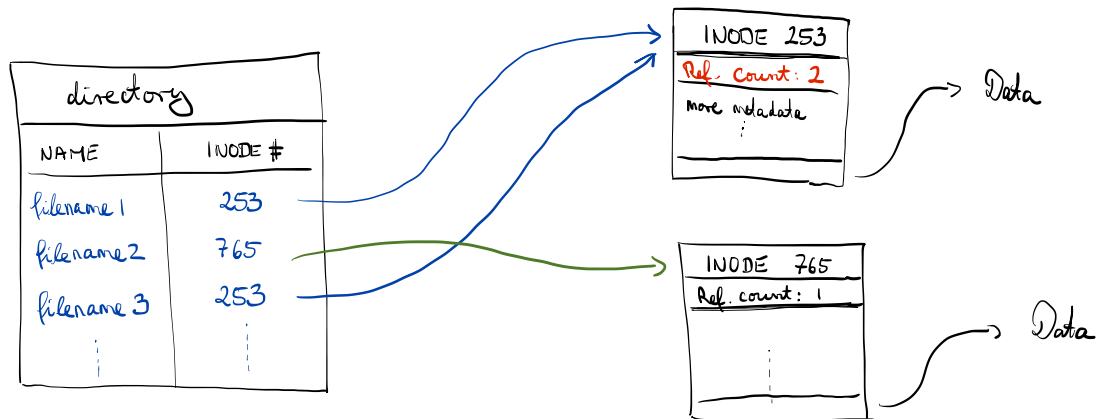
SMB/CIFS (comes from Windows): complex, has access authentication

## File system links

Purpose: provide access to files via alternative names, without moving or copying file data.

## Option 1: "Hard links"

Background: a directory is a collection of inode numbers pointing to the actual file data. Thus, it is natural to allow several entries to point to the same inode. (Thus, every file entry in a directory could be seen as a hard link, the point is that there may be more than one such link, violating the tree structure.)



- File data gets deleted when ref. count drops to zero.
- Can only create hard links on the same FS partition (NOT across mounts!)
- Hardlinking to directories is dangerous (might create cycles), forbidden on some OSs.

### Advantages:

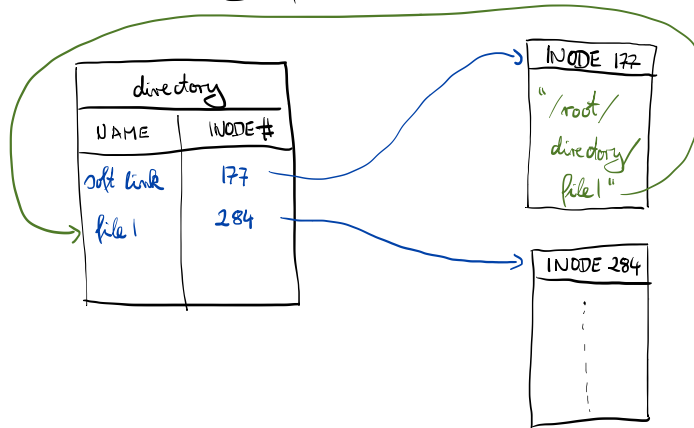
- natural concept (use existing pointer structure)
- no overhead
- can be used to save disk space (e.g. 'git clone' on same FS partition)

### Concerns:

- backup tools and disk usage monitors need to be aware
- can cause unexpected behavior for user-facing programs → see exercises!

## Option 2: "soft links" or "symbolic links"

A soft link is a string representation of the target path:



- following a soft link requires one additional path resolution  
→ slower than hard link
- deleting the link will never delete the target!
- A soft link is visibly distinct from a hard link to the same file

### Advantage:

- May cross FS-boundaries (i.e., soft links into a mounted FS work)

## RAID ("Redundant Array of Inexpensive Disks")

Idea: Use many cheap standard disk drives to increase reliability and/or speed of access

The following operating modes are important (see web links for details):

RAID 0 ("striping"): Spread data across several drives  
→ increases speed, decreases reliability

RAID 1 ("mirroring"): Mirror data onto two or more drives  
→ increases read speed  
→ write speed is speed of slowest drive  
→ improve reliability  
→ no increase in capacity

RAID 5 ("striping with parity"): Like striping, but with a parity block that stores the parity of each bit position across all drives

- Like RAID 0 plus one disk (but in practice, parity data is distributed across all drives to avoid write bottleneck)
- Significant improvement in reliability
- Significant improvement in speed
- Capacity of an  $n$ -disk array is joint capacity of  $n-1$  disks

RAID 6: Like RAID 5, but with additional redundancy to protect against failure of two disks. For large-capacity disks where the time to restore a single failed disk can be significant, increasing the chances for a second-disk failure, thus complete loss of data under RAID 5.