

Media and git makeup (and kill svn)

A typical media repository

- Very large individual files, 100mb
- Very large complete repository 100's gbs
- Does not delta compress well

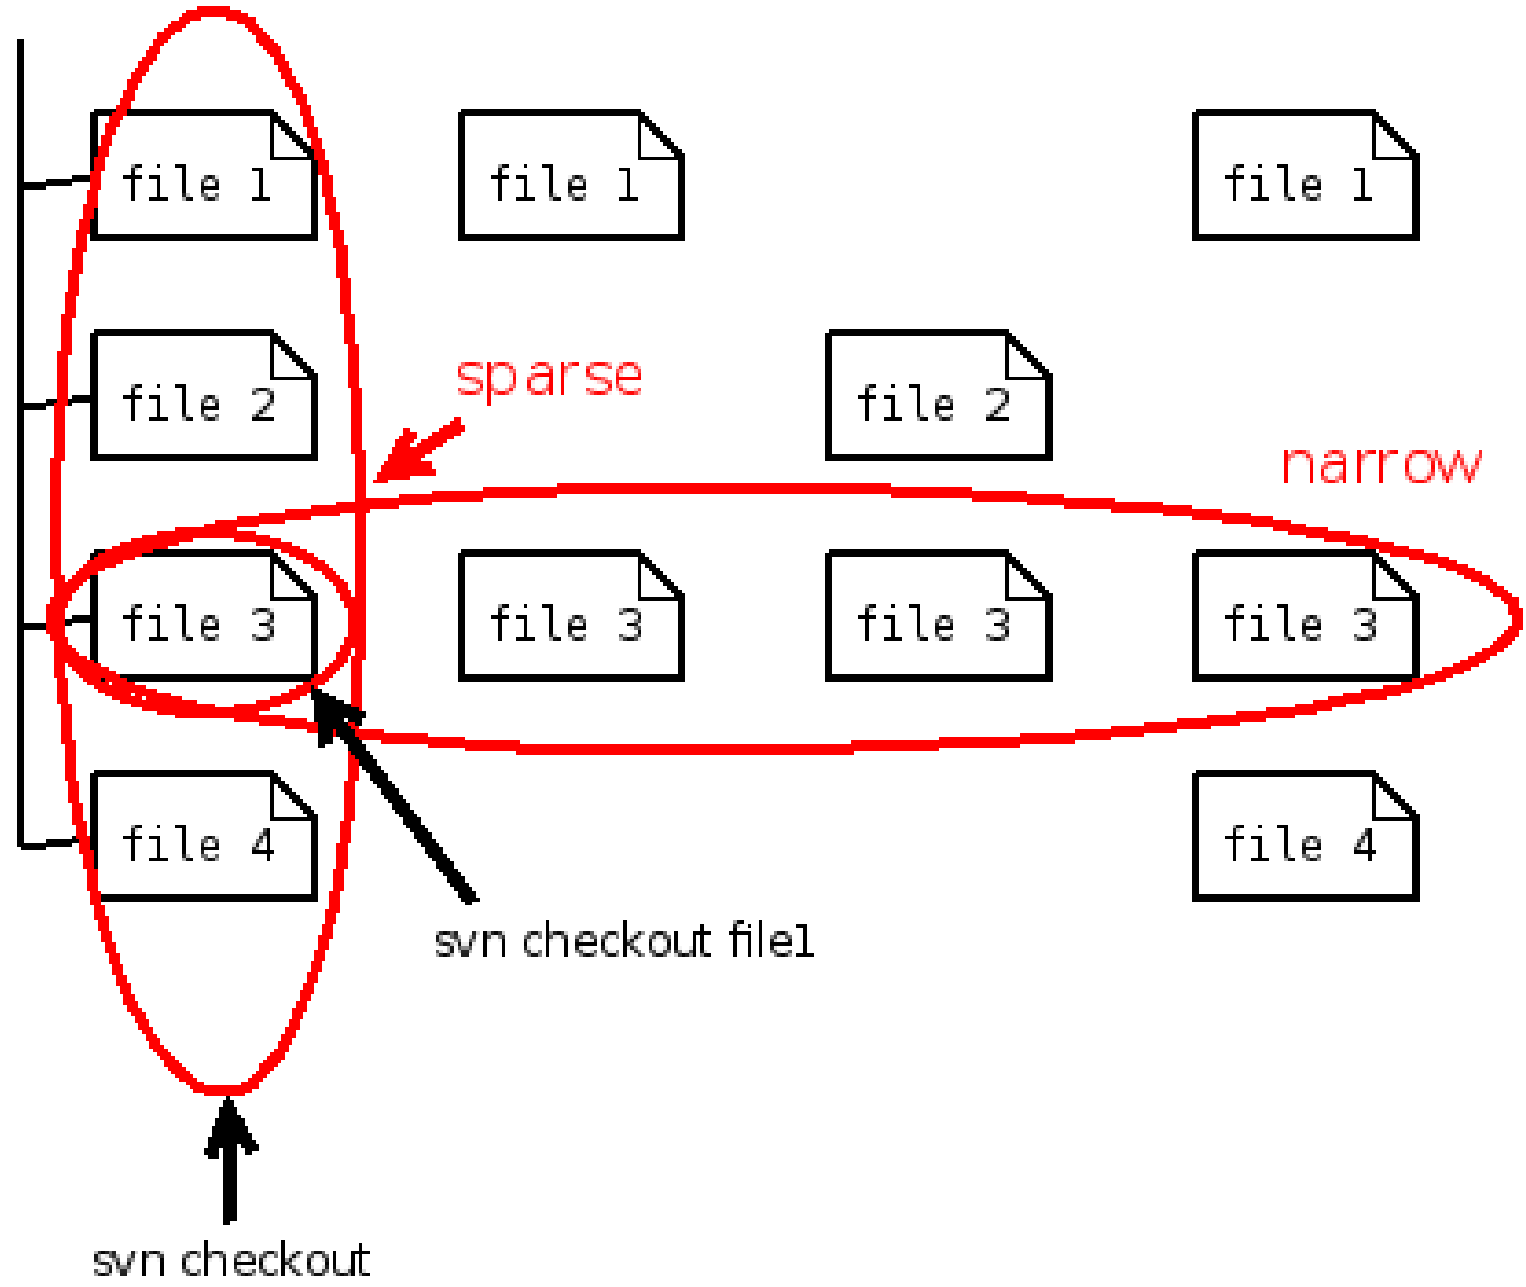
r10

r9

r8

r7

shallow



Current git problems

- Clone always gets complete history (100gb)
- Shallow commits are second-class citizens
(can not commit)
- Can not do narrow checkouts

The solution

- Always clone the full metadata.
 - Metadata is trivial in size < 100mb.
 - Metadata is the most interesting.
- Only clone the blobs you are interested in.

- Commits work!
- Commit log works!
- Only slow as svn in worst case.

Making it work

- git-daemon changes
 - Download non-blobs
 - Download individual blobs
- git-clone changes
 - Record which blobs exist but not downloaded
 - Needed for git-fsck, git-gc, etc
- git-commit, support ignoring non-existent
- git tools
 - Understand the difference between "non-existence" and "don't have"

Problems

- Merging "sideways"
 - If you don't have a blob that needs to be merged it's impossible to merge.
- Automatic download vs Manual downloads
 - Automatic means "shoot yourself in foot"
 - Manual makes operations a pain