

## Aggressive Filesystem Caching in Plan 9

*Geoff Collyer*

geoff@plan9.bell-labs.com

Bell Laboratories

Murray Hill, New Jersey 07974

USA

### ABSTRACT

As part of scaling Plan 9 to run on very large machines, such as IBM's Blue Genes, we needed to reduce the load on the Plan 9 file server(s) used by a Blue Gene Plan 9 cluster. Plan 9 already had a caching file server, *cfs*, but it still passed a lot of 9P traffic through to the file server it was caching, which is usually unnecessary for our Blue Gene environment, in which the files being served are largely static. *Ramcfs* and *fscfs* provide more aggressive caching than *cfs* and greatly reduce 9P traffic passed through.

### Ramcfs

When many instances of Plan 9 start running at the same time, they attempt to read the same set of files at the same time. This was flooding our file server with requests for traffic, much of which was duplicated. *Ramcfs* was our first attempt to avoid flooding our file server by caching files in memory. In many cases, *ramcfs* can answer 9P requests immediately, from in-memory data, without having to communicate with the file server at all. In other cases, it has to check file version numbers at least.

### Changes to create ramcfs

*Ramcfs* was created by making a copy of *cfs* and changing it to use a region of memory as its cache rather than a region of disk. *Ramcfs* always 'formats' the cache when it starts, since it does not persist from run to run. Further improvements were to treat files with `qid.ver`s == 0 as synthetic thus uncachable and to add a `-r` option that declares the files being cached to be unchanging (readonly) and thus capable of being cached more aggressively.