It was not being used recursively, and there are no subdirectories in
the directory being read.

David - Prayer is probably the correct solution, though not nearly as
satisfying as a good rant or monitor punch.

I conducted a test by performing a CD into the desired Samba-connected
directory and then performed FILE_SEARCH without a directory specified,
but it was not any faster than specifying the directory within
FILE_SEARCH itself.

It appears there is some dependency on the number of filenames
returned. If I search the same directory but specify a wildcard that
returns fewer files, the search time is shorter. # of files in the
directory was the same for each test.

IDL> print,systime() & z=file_search('20060119*.nogaps') &
print,systime()
Thu Jan 19 10:49:06 2006
Thu Jan 19 10:49:51 2006
IDL> help,z
Z   STRING   = Array[952]
(45 seconds to return 952 filenames)

IDL> print,systime() & z=file_search('2006011906*.nogaps') &
print,systime()
Thu Jan 19 10:51:15 2006
Thu Jan 19 10:51:24 2006
IDL> help,z
Z   STRING   = Array[136]
(9 seconds to return 136 filenames)

I created a directory with 2000 files on the local hard drive and both
FILE_SEARCH and FINDFILE read them within a second, so the sheer number
of filenames is not the issue. So my guess is that FILE_SEARCH and
Samba don't play efficiently together.

- Mark