Subject: Re: FINDFILE vs. FILE_SEARCH

Posted by David Fanning on Wed, 18 Jan 2006 16:29:25 GMT

View Forum Message <> Reply to Message

Mark Conner writes:

- > Workarounds (other than continuing to use FINDFILE and hope it works
- > properly)?

I'm a strong proponent of prayer in software development. It is always more effective than strong language. :-)

Cheers.

David

--

David Fanning, Ph.D.

Fanning Software Consulting, Inc.

Coyote's Guide to IDL Programming: http://www.dfanning.com/

Subject: Re: FINDFILE vs. FILE_SEARCH

Posted by btt on Wed, 18 Jan 2006 16:52:06 GMT

View Forum Message <> Reply to Message

Mark Conner wrote:

- > I'm using IDL 6.1 on WinXP Pro to look for a set of files (can be
- > 1500+) in a directory on a Sun system connected via Samba. The
- > directory appears as the standard UNC (\\server\\directory\\files.ext).

>

- > I've read all the stuff that says to use FILE_SEARCH instead of
- > FINDFILE. However, it takes nearly two minutes to return a list of
- > files (setting /NOSORT saves perhaps a couple seconds). Using FINDFILE
- > I get a list back in three seconds. Any ideas why?

>

- > Network traffic or machine loads on either end are not a factor. I'm
- > guessing it's a issue with how FILE_SEARCH interacts with
- > Samba-connected directories, but does anyone know for sure?
- > Workarounds (other than continuing to use FINDFILE and hope it works
- > properly)?

>

Are you using FILE_SEARCH in a recursive mode? I find that result = FILE_SEARCH("path", "name") can take a lot longer than result = FILE_SEARCH("path/name") if there are many subdirectories to search.

Subject: Re: FINDFILE vs. FILE_SEARCH
Posted by news.qwest.net on Wed, 18 Jan 2006 17:26:09 GMT
View Forum Message <> Reply to Message

"Mark Conner" <mconner1@gmail.com> wrote in message news:1137600968.365781.139820@z14g2000cwz.googlegroups.com...

- > I'm using IDL 6.1 on WinXP Pro to look for a set of files (can be
- > 1500+) in a directory on a Sun system connected via Samba. The
- > directory appears as the standard UNC (\\server\directory\files.ext).

>

- > I've read all the stuff that says to use FILE_SEARCH instead of
- > FINDFILE. However, it takes nearly two minutes to return a list of
- > files (setting /NOSORT saves perhaps a couple seconds). Using FINDFILE
- > I get a list back in three seconds. Any ideas why?

>

- > Network traffic or machine loads on either end are not a factor. I'm
- > guessing it's a issue with how FILE_SEARCH interacts with
- > Samba-connected directories, but does anyone know for sure?
- > Workarounds (other than continuing to use FINDFILE and hope it works
- > properly)?

>

> - Mark

In the past I have run into a lot of problems with the various filename reading functions.

Some suggestions:

- use findfile (if it does indeed give a proper list of files)
- spawn to the system and get a file listing
- perhaps change the current directory (IDL> cd,'current_dir') then do the filesearch function (then change the cd back)
- create a listing in the operating (>ls > dir.txt) then read that text file into IDL

(if the file listing is fairly static, this should work fine)

- create a rule for predicting what filenames should exist (for instance, with data

files with names like "wind20060101.dat', you can just create the filename, then see if it exists).

Cheers,

bob

Subject: Re: FINDFILE vs. FILE_SEARCH
Posted by Mark Conner on Thu, 19 Jan 2006 17:06:06 GMT
View Forum Message <> Reply to Message

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