Subject: Re: big file navigation
Posted by R.G. Stockwell on Tue, 28 Oct 2003 19:53:09 GMT
View Forum Message <> Reply to Message

"Ben Tupper" <btupper@bigelow.org> wrote in message
news:bnm487$1317fk$1@ID-189398.news.uni-berlin.de...
...
> IDL> for i = 0L, 55855-1 do readU, U,d
>
>
> Problem solved! I can break the file into smaller 'chunks' of 4GB by
reading in
> images and writing them to a new file. OK!
>
> But here's what I wish I understood (and could exploit!) Somewhere the
system is
> keeping track of the file location so that READU knows where to read. So,
deep
> down in its guts, IDL is navigating its way through with 64-bit file
offsets.
> Ain't it? Is there anything here I could get my hooks into so I don't
have to
> break the original file into smaller ones?
>
> Cheers,
> Ben

The point_lun seems to use a 64 bit offset

from the help:

If Unit is negative, Position must be a named variable into which the
current file position will be stored. The returned type will be a longword
signed integer if the position is small enough to fit, and an unsigned
64-bit integer otherwise.

Cheers,
bob

Subject: Re: big file navigation
Posted by btt on Tue, 28 Oct 2003 21:33:24 GMT
View Forum Message <> Reply to Message

R.G. Stockwell wrote:
> "Ben Tupper" <btupper@bigelow.org> wrote in message
> news:bnm487$1317fk$1@ID-189398.news.uni-berlin.de...

> ...
>
>> IDL> for i = 0L, 55855-1 do readU, U,d
>>
>>
>> Problem solved! I can break the file into smaller 'chunks' of 4GB by
>
> reading in
>
>> images and writing them to a new file. OK!
>>
>> But here's what I wish I understood (and could exploit!) Somewhere the
>
> system is
>
>> keeping track of the file location so that READU knows where to read. So,
>
> deep
>
>> down in its guts, IDL is navigating its way through with 64-bit file
>
> offsets.
>
>> Ain't it? Is there anything here I could get my hooks into so I don't
>
> have to
>
>> break the original file into smaller ones?
>>
>> Cheers,
>> Ben
>
>
> The point_lun seems to use a 64 bit offset
>
> from the help:
>
> If Unit is negative, Position must be a named variable into which the
> current file position will be stored. The returned type will be a longword
> signed integer if the position is small enough to fit, and an unsigned
> 64-bit integer otherwise.
>
>
> Cheers,
> bob
>
>
Hi Bob,

I saw that and thought it was quite interesting. So, I can find out where I am
within a file - but I can't go anywhere within it.

Ben

Subject: Re: big file navigation
Posted by R.G. Stockwell on Tue, 28 Oct 2003 21:42:37 GMT
View Forum Message <> Reply to Message

"Ben Tupper" <btupper@bigelow.org> wrote in message
news:bnmnb6$12g7c2$1@ID-189398.news.uni-berlin.de...
> Hi Bob,
>
> I saw that and thought it was quite interesting. So, I can find out where
I am
> within a file - but I can't go anywhere within it.
>
> Ben

Point lun will point the lun wherever you want when used with a positive lun
eg. point_lun,1,1024
Point lun will read the current position when used with a negative lun
eg. point_lun,-1,curpos

I don't have a file that is > 4gb, so I don't know if you are running into
bugs there, but this does work fine for "normal" files.

So you can write you own routine to do what associate does.
(actually I am surprised these things fail when the file is > 4gb)

Cheers,
bob

Subject: Re: big file navigation
Posted by btt on Fri, 31 Oct 2003 14:00:39 GMT
View Forum Message <> Reply to Message

R.G. Stockwell wrote:
> "Ben Tupper" <btupper@bigelow.org> wrote in message
> news:bnmnb6$12g7c2$1@ID-189398.news.uni-berlin.de...
>
>> Hi Bob,
>>

>> I saw that and thought it was quite interesting. So, I can find out where
>
> I am
>
>> within a file - but I can't go anywhere within it.
>>
>> Ben
>
>
>
> Point lun will point the lun wherever you want when used with a positive lun
> eg. point_lun,1,1024
> Point lun will read the current position when used with a negative lun
> eg. point_lun,-1,curpos
>
> I don't have a file that is > 4gb, so I don't know if you are running into
> bugs there, but this does work fine for "normal" files.
>
> So you can write you own routine to do what associate does.
> (actually I am surprised these things fail when the file is > 4gb)
>

Hi again,

What you describe above does work for "normal" sized files, but it doesn't for the "wicked big files" I have in hand. I tried to pull together something like ASSOC - but the limiting code was POINT_LUN.

My example using the 8+GB file with output is below.

Thanks for the ideas, but I think I'll have to break the file into smaller pieces.

Cheers,
Ben

******* Code Starts here ********

```
PRO Big_File_Test

file = '/Users/Shared/data/VPR16/vpr16.stk'

Help, File_Info(file), /str

openR, U, file, /GET_LUN

location = 5000000LL ; five million as starting point

increment = 10LL
```

```
 ;simply loop through, incrementing the requested location
 ;by a factor of 10 each time.
While NOT EOF(U) Do Begin
 Print, "Requested location in bytes is...", location
 Point_LUN, U, location
 location = location * increment
EndWhile

Free_LUN, U
END ;big_File_Test
```

******** Code End here ********

++++++++++ Output begins here

```
IDL> big_file_test
** Structure FILE_INFO, 21 tags, length=64, data length=63:
   NAME            STRING    '/Users/Shared/data/VPR16/vpr16.stk'
   EXISTS          BYTE         1
   READ            BYTE         1
   WRITE           BYTE         1
   EXECUTE         BYTE         0
   REGULAR         BYTE         1
   DIRECTORY       BYTE         0
   BLOCK_SPECIAL   BYTE         0
   CHARACTER_SPECIAL
                   BYTE         0
   NAMED_PIPE      BYTE         0
   SETUID          BYTE         0
   SETGID          BYTE         0
   SOCKET          BYTE         0
   STICKY_BIT      BYTE         0
   SYMLINK         BYTE         0
   DANGLING_SYMLINK
                   BYTE         0
   MODE            LONG             420
   ATIME           LONG64        1067607855
   CTIME           LONG64        1067297307
   MTIME           LONG64        1067297307
   SIZE            LONG64          -10606592
Requested location in bytes is...           5000000
Requested location in bytes is...          50000000
Requested location in bytes is...         500000000
Requested location in bytes is...        5000000000
Requested location in bytes is...       50000000000
```

```
% POINT_LUN: IDL on this platform does not support 64-bit file access: LOCATION.
% Execution halted at: BIG_FILE_TEST      15
/Users/ben/pemaquid/lib/grampus/big_file_test.pro
%                    $MAIN$

---------- Output ends here
```