Subject: Re: Readu, Writeu Causing Segmentation Fault
Posted by A.R. on Wed, 09 Sep 2009 21:41:58 GMT
View Forum Message <> Reply to Message

On Sep 7, 1:10 am, Maarten <maarten.sn...@knmi.nl> wrote:
> On Sep 7, 6:21 am, "A.R." <alrom...@gmail.com> wrote:
>
>> I have created rather large volumes (3-D binary arrays) in IDL that I
>> save as unformatted data (i.e. writeu). When the volume gets above a
>> certain size (around 2.2 GB) when I attempt to writeu or readu the
>> volume, I am kicked out of IDL with a segmentation fault. I don't
>> think it's a memory issue, I'm running 64-bit IDL on a 64-bit linux
>> running redhat enterprise with 12 GB of ram.
>
> Could it be a file-system issue? Or perhaps IDL uses 32-bit integers
> somewhere for the internal file pointer, you never know...
>
>> One workaround I've used for 'writeu' is to write to the file in
>> chunks, by using a for loop to write each slice of the volume
>> individually. This works for saving the volume, but then I have no
>> way to re-open the volume in IDL without the readu segmentation fault!
>
>> Anyone have any ideas for what could be causing this? I'm running IDL
>> Version 6.2.
>
>> I appreciate any suggestions/ideas you smart people might have! In
>> the meantime, I'll continue banging my head against the wall over this
>> one.
>
> IDL 6.2 can just do this: useHDF-5 as an alternative to writeu/readu.
> It will allow you to selectively write and read chunks of data,
> without size limitations (note before you run out of 12 GB main
> memory).HDF-4 won't let you write file of the size you apparently
> need to use, neither will netcdf.
>
> Be aware though that IDL 6.2 is the first version with write support
> forHDF-5, and if I recall correctly, you may encounter some funny
> features that have since been removed...
>
> Best,
>
> Maarten

Hello,

I have attempted to use HDF-5, and can successfully write it to the
file, however when I re-open the file it loses its dimensions, i.e. :
Writing the HDF-5 File

```
IDL> test=bytarr(1320,1250,1414)
IDL> fileID=H5F_CREATE('testfile.h5')
% Loaded DLM: HDF5.
IDL> datatypeID=H5T_IDL_CREATE(test)
IDL> dataspaceID=H5S_CREATE_SIMPLE(2333100000)           ;  Test
volume originally 1320x1250x1414
IDL> datasetID=H5D_CREATE(fileID, 'testdata',datatypeID,dataspaceID)
IDL> H5D_WRITE,datasetID,test
IDL> H5F_CLOSE,fileID
IDL> exit

**Trying to re-open HDF-5 File**
IDL> vol=H5_PARSE('testfile.h5', /READ_DATA)
% Compiled module: H5_PARSE.
% Loaded DLM: HDF5.
IDL> print, size(vol)
     1     1     8     1
IDL> exit
```

My question: How do I get the correct dimensions back using HDF-5? I'm guessing it's one of the settings I should supply when writing to the file, but I'm new to HDF-5 . . .

In response to Wox (post below) I can also writeu that file, but for some reason when I have a large number of actual ELEMENTS in the volume (apparently not file size) is when I get the segmentation fault. And the magic number seems to be ~2,300,000,000 (~2^32) . . . coincidence? Doubtful. But how do I fix it?

Thanks!