Subject: Re: IDL 8.2.2 released
Posted by chris_torrence@NOSPAM on Fri, 08 Feb 2013 16:14:53 GMT
View Forum Message <> Reply to Message

On Thursday, February 7, 2013 9:47:13 PM UTC-7, timoth...@gmail.com wrote:
>> Also, if you have the CR number for a particular bug, I'd be happy to check on its status for you.
>
>>
>
>>
>
>>
>
>> mp
>
>
>
> Hi Mark,
>
>
>
> I would like to take you up on your offer. CRS 67311
>
> Its about the very slow restore times of large arrays from .sav files. We are using IDL in an emergency warning system where every second counts.
>
>
>
> Thanks,
>
> Tim

Hi Tim,

This bug is on the plate for IDL 8.3. I just did some quick tests, and it looks like some of the speed difference is because the IDL save files are saved in "big endian" format. When you read this on a Windows, Linux, or Mac machine, it needs to convert all of the arrays to little endian. Something we are doing in the C code must be slower than the Python big->little endian conversion.

In the meantime, if you are desperate for a speedup, you could use just an unstructured readu/writeu. It's not as flexible as the save file, but it will be as fast (if not faster) than the Python. On my Mac, using your original reproduce code, I could read that 400 Mb array in 6 seconds using a save file, and 0.6 seconds using a straight readu.

Just out of curiosity, are you using scipy's "idlsave" package to read IDL save files into Python?

-Chris