
Subject: Re: lost data?

Posted by [George N. White III](#) on Thu, 10 May 2001 13:24:13 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Thu, 10 May 2001, src wrote:

> Is there a bug in IDL's Save/Restore command? I've just spent the last 18
> days running a Monte Carlo simulation to seemingly lose all my data. The
> problem occurred when our license manager stopped responding (network
> problem) hence the IDL session running the simulation crashed. The MC
> code is designed to save results periodically as it runs (just in case
> this sort of thing happens). I've just tried:
>
> "restore, 'mc_file.sav', /Verbose"
>
> only to get:
>
> % RESTORE: IDL version 5.3 (linux, x86).
> % RESTORE: Truncated save file, restored as much as possible:
>
> That "restored as much as possible:" is in fact 0 (zero). Despite the file
> itself being 17 Mb! Some of my .sav files are a lot bigger than this, yet
> don't seem to have any problems. Is there anyway to recover this file, or
> prevent this happening again in the future? I'm going to very upset to
> lose 18 days work...
>
> cheers,
> S

Some OS's (SGI Irix) have a checkpoint facility that works at the level of processes, and doesn't require support built in to the application. I know there has been some work on checkpointing for linux, as the same capability is required to migrate a process to a new node in some distributed processing systems. I don't know if there is anything you can use with IDL, but it is certainly worth a look.

We run IDL batch jobs on a compute server that almost never goes down (big UPS and generator), but some jobs want an X-server, so the users have been setting the DISPLAY variable to an X-server on a workstation that doesn't have generator power. The jobs die if the power is out too long for the UPS's that run the network and workstations.

The trouble with the things you do to try to improve reliability for long batch runs is that it is almost impossible to test all the things that can go wrong -- power failures, disks getting full or failing, network failures, etc. Do other people have similar cautionary tales? What changes were needed to make batch processing more robust?

--

George N. White III <gnw3@acm.org> Bedford Institute of Oceanography
