
Subject: Re: Multiple journal files on linux
Posted by [Ken Mankoff](#) on Wed, 19 Jun 2002 15:57:55 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 19 Jun 2002, Mark McGrath wrote:

- > This routine makes an `almost' unique file. An example:
- > Running IDL on a linux cluster, as we do here, if the machine clocks
- > are a little out of sync it would be possible to over write the file
- > that had been opened already, no? (Not much info lost there though,
- > but the more serious case of the clocks being a lot out of sync then
- > large amounts of info in the journal file could be lost.)
- >
- > Any ideas, anyone?

I ran into the problem of making a unique directory on a multiprocessor machine, where a timestamp may not suffice...

Use the unique process ID. Each IDL session will have a unique PID, no matter how many processors and what the time is on each machine... I use a combination, the time (in msec), then a "_", then the PID.

I believe you can still have multiple processes assigned the same time and PID if you are using networked computers (as opposed to 1 computer with multiple processors), but the chances, especially when using msec and not sec, drop to VERY low.

-k.