

art - Bug #6960

sqlite error when running nova calibration subrun summation job in nova development release

09/09/2014 05:16 PM - Alexander Radovic

Status:	Closed	Start date:	09/09/2014
Priority:	Immediate	Due date:	
Assignee:	Christopher Green	% Done:	100%
Category:	I/O	Estimated time:	0.00 hour
Target version:	1.12.00	Spent time:	4.00 hours
Occurs In:		Experiment:	NOvA
Scope:	Internal	SSI Package:	art

Description

Hey Artists,

Chris Green has asked me to post here a recent error which I saw whilst running a simple job in the new version of art over files produced with the old version of art. The original email:0

So I checked out a new development release and went to run sumsubrunsjob.fcl in the hopes of checking that the new version of art does in fact fix the calibration memory woes.

However I got the following new error:

```
<novagpvm05.fnal.gov> nova -c sumsubrunscalibjob.fcl /nova/ana/users/radovic/exampleAttenPC/*
attenprof*root
%MSG-i MF_INIT_OK: nova 09-Sep-2014 15:36:44 CDT JobSetup
Messagelogger initialization complete.
%MSG
09-Sep-2014 15:36:48 CDT Initiating request to open file /nova/ana/users/radovic/exampleAtten
PC/fardet_r00014702_s02_t02_S14-06-09_v1_data.attenprof.root
09-Sep-2014 15:36:48 CDT Successfully opened file /nova/ana/users/radovic/exampleAttenPC/fard
et_r00014702_s02_t02_S14-06-09_v1_data.attenprof.root
%MSG-s ArtException: nova 09-Sep-2014 15:36:48 CDT JobSetup
cet::exception caught in art
---- Configuration BEGIN
  FailedInputSource Configuration of main input source has failed
---- SQL error BEGIN
  SQLite error: database disk image is malformed (11): database disk image is malformed
---- SQL error END
---- Configuration END
%MSG
```

Art has completed and will exit with status 9.

has anyone else seen an SQLite error in development?

I've just checked and I don't see this error in S14-08-19.

Some additional information:

The location of any files suffering from the problem (all of the files seem to suffer from this problem):

/nova/ana/users/radovic/exampleAttenPC/*attenprof*root

The NOvA release designation and the output of the ups active | sort command for the releases that:

generated the file (attached S14-06-09upsact.txt);

successfully read the file (attached S14-08-19upsact.txt);

failed to read the file (attached devupsact.txt).

The file config.fcl which should be the result of the command:

ART_DEBUG_CONFIG=config.fcl nova -c sumsubrunscalibjob.fcl /nova/ana/users/radovic/exampleAttenPC/*attenprof*root
(attached)

Regenerating the files with devupsact.txt creates files which devupsact.txt can read.

cheers,
-Alexander

Associated revisions

Revision 914902e9 - 09/15/2014 02:02 PM - Christopher Green

Fix for bug which gave rise to issue #6960.

History

#1 - 09/12/2014 11:46 AM - Christopher Green

- Description updated
- Category set to I/O
- Status changed from New to Assigned
- Assignee set to Christopher Green
- Priority changed from Normal to Immediate
- SSI Package fhicl-cpp added
- SSI Package deleted ()

Could you tell me whether the error is also seen with config_dumper?

#2 - 09/12/2014 04:53 PM - Christopher Green

- Status changed from Assigned to Resolved
- Target version set to 1.12.00
- % Done changed from 0 to 100
- SSI Package art added
- SSI Package deleted (fhicl-cpp)

I have reproduced the behavior.

The proximal problem is that art is now better at detecting errors returned by calls to SQLite. These files are in fact malformed due to a well-intentioned attempt to truncate the parameter set database embedded in the ROOT file using, if I remember correctly, a small program written by Chris Backhouse based on config_dumper. The root cause reason why this resulted in a malformed ROOT file in the first place was a minor bug in art's "tkeyvfs" SQLite extension responsible for the reading and writing of SQLite data from and to ROOT files, which caused the duplication of keys rather than a new cycle of the existing key. I should point out that this cannot happen in a normal art execution, but is entirely an artifact of the after-the-fact file editing.

The older versions of art still failed to open the database but went blithely on their way -- incidentally having the intended effect, that of faster file open times.

Going forward, the problem this file-edit was intended to solve has been solved by other means (see issue [#5805](#) and [Release Notes 1.11.00](#)). Files currently suffering from the TKey duplication problem may be fixed with the following heuristic:

1. Open each file in ROOT, in update mode.
2. Get the list of keys from the file.
3. For each TKey (including different cycles), check the size. One of them will be much larger than the others -- this is the one you keep. Remove the others by calling the Delete() method on the TKey.
4. Close the file.

The problem which gave rise to the multiple keys has been fixed with commit:f658e44.

#3 - 11/13/2014 02:59 PM - Christopher Green

- Status changed from Resolved to Closed

Files

config.fcl	4.35 KB	09/09/2014	Alexander Radovic
devupsact.txt	6.74 KB	09/09/2014	Alexander Radovic
S14-06-09upsact.txt	5.93 KB	09/09/2014	Alexander Radovic
S14-08-19upsact.txt	6.14 KB	09/09/2014	Alexander Radovic