

art - Bug #21316

v2_11 doesn't fix provenance from v2_10

11/06/2018 11:45 AM - Raymond Culbertson

Status:	Closed	Start date:	11/06/2018
Priority:	Normal	Due date:	
Assignee:	Kyle Knoepfel	% Done:	100%
Category:	Infrastructure	Estimated time:	4.00 hours
Target version:	2.11.05	Spent time:	0.00 hour
Occurs In:	2.11.04	Experiment:	Mu2e
Scope:	Internal	SSI Package:	art

Description

In this issue

<https://cdcv.sfnal.gov/redmine/issues/19465>

a provenance problem was found in art v2_10 and fixed in v2_11.

Now we are testing art v2_11_01 and still see a form

of the problem. However we are not sure if the fact that

the input files were created in v2_10 means that they will

always be corrupt, or they should be read cleanly by v2_11.

Below is the how to reproduce the current error. Is this

expected? Is there a way to repair files or suppress the error?

In process 1:

```
source /cvmfs/mu2e.opensciencegrid.org/setupmu2e-art.sh
```

```
source /cvmfs/mu2e.opensciencegrid.org/Offline/v7_1_2/SLF6/prof/Offline/setup.sh
```

```
(sets up art v2_10_04)
```

```
export
```

```
IN=/pnfs/mu2e/tape/phy-sim/dig/mu2e/CeEndpoint/MDC2018b/art/13/a3/dig.mu2e.CeEndpoint.MDC2018b.001002_00000000.art
```

```
mu2e -n 10 -c JobConfig/common/artcat.fcl -s $IN -o copy_old.art
```

In process 2

```
source /cvmfs/mu2e.opensciencegrid.org/setupmu2e-art.sh
```

```
source /cvmfs/mu2e.opensciencegrid.org/Offline/v7_1_2a/SLF6/prof/Offline/setup.sh
```

```
(sets up art v2_11_01)
```

```
mu2e -s copy_old.art -c Print/fcl/print.fcl
```

```
produces error: Product reported as not present, but is pointed to nonetheless!
```

Note that if the steps in process 1 are done in process2, creating

copy_new.art, the result will be read cleanly by v2_11. If copy_old.art

is read and written by v2_11, then the provenance is read in v2_11, it

produces the error, so v2_11 isn't repairing the bad file. Is this expected?

The releases can be created locally by

```
git clone http://cdcv.sfnal.gov/projects/mu2eoffline/mu2eoffline/Offline.git
```

```
git checkout v7_1_2 (or v7_1_2a)
```

```
source /cvmfs/mu2e.opensciencegrid.org/setupmu2e-art.sh
```

```
source Offline/setup.sh (instead of cvmfs)
```

```
scons -j 20
```

History

#1 - 11/06/2018 12:10 PM - Kyle Knoepfel

- Status changed from New to Feedback

Please update to [art2.11.03](#) to determine if there is still a problem. Issue [#19809](#) addresses a similar issue that was resolved for [art2.11.02](#).

#2 - 11/14/2018 03:21 PM - Kyle Knoepfel

From Ray:

This problem as described in the original report still persists when use art v2_11_04 as the "process 2" art version.

#3 - 11/14/2018 03:52 PM - Kyle Knoepfel

- Status changed from *Feedback* to *Assigned*
- Assignee set to *Kyle Knoepfel*

Investigating.

#4 - 11/14/2018 03:52 PM - Kyle Knoepfel

- Occurs In *2.11.04* added
- Occurs In *deleted (2.11.01)*

#5 - 11/15/2018 09:27 AM - Kyle Knoepfel

- Category set to *Infrastructure*
- Status changed from *Assigned* to *Feedback*
- % Done changed from *0* to *100*
- Estimated time set to *4.00 h*

The problem is understood. Although the changes made for [art2.11.02](#) addressed the persistency issues when using the RootOutput module, the workflow here uses a ProvenanceDumper-based module, which is independent of RootOutput and was not updated for [art2.11.02](#). In addition there were some mistaken assumptions regarding how *art* calculates the availability of the product.

The following commits resolve these problems:

- [canvas:01adf4d](#)
- [art:517272b](#)

Does Mu2e request a bug-fix release (e.g. 2.11.05), or shall we just include the fix for the next planned release of *art*, which is version [3.02.00](#)? I suspect the former.

#6 - 11/16/2018 01:06 PM - Kyle Knoepfel

- Target version set to *2.11.05*

#7 - 11/16/2018 01:45 PM - Kyle Knoepfel

- Status changed from *Feedback* to *Resolved*

Per discussion with Rob, the fix for this bug will be included with [art2.11.05](#).

#8 - 11/19/2018 03:01 PM - Kyle Knoepfel

- Status changed from *Resolved* to *Closed*