

art - Bug #6823

Metadata field "parents" is not compatible with SAM

08/18/2014 04:38 PM - Nathan Mayer

Status:	Closed	Start date:	08/18/2014
Priority:	Immediate	Due date:	
Assignee:	Christopher Green	% Done:	100%
Category:	Metadata	Estimated time:	2.00 hours
Target version:	1.11.00	Spent time:	4.00 hours
Occurs In:		Experiment:	NOvA
Scope:	Internal	SSI Package:	art

Description

In art v1_10_00b the SAM metadata parent field is being set by art and nothing else has been able to override it. Further more the parentage is being set to the full path to the input file. This is incompatible with SAM which just expects a file name (no path).

Under certain circumstances we need to be able to override this and set the parentage ourselves. And we are currently unable to do this either.

Associated revisions

Revision 8ed95f9f - 08/19/2014 02:23 PM - Christopher Green

Issue #6823.

History

#1 - 08/19/2014 02:32 PM - Christopher Green

- Subject changed from Metadata field "parent" is not compatible with SAM to Metadata field "parents" is not compatible with SAM
- Status changed from New to Resolved
- Assignee set to Christopher Green
- Priority changed from High to Immediate
- Target version set to 1.11.00
- % Done changed from 0 to 100
- Estimated time set to 2.00 h
- SSI Package art added
- SSI Package deleted ()

Fixed with [8ed95f9](#).

- The file names stored in the parents SAM metadata item are now basename-only.
- The FileCatalogMetadata service now takes a parameter wantFileParentsMetadata, which defaults to true. Each instance of RootOutputFile also takes a parameter of the same name whose value (if set) overrides that of the service for that stream only. If this is set to false, the file parentage metadata may be set explicitly (following JSON syntax for sequences) either globally by calling FileCatalogMetadata::addMetadata(...) or locally from a FileCatalogMetadataPlugin for individual output streams (see [FileCatalogMetadataPlugin](#) for more details).

Does this resolve your SAM issues?

#2 - 08/19/2014 03:24 PM - Gavin Davies

Chris, that's great! This was a big issue for us in the current ART version.

Reading your fix description it sounds like this will allow us to specify the fcl file as the parent, over-riding the parents field that is currently grabbing the intermediate file that we throw away. If we can force it to be the fcl file with this fix (as we could previously) then this is exactly what we need.

Adam A. listed a few other issues that have come up with SAM+ART interfacing recently (although he's on holiday this week). This is what he said:

- 1) file_format_era and file_format_version do not exist as parameters in SAM. We have had to remove these using our metadata extractor to get our files accepted by SAM.
- 2) first_event and last_event are expected to be integers by SAM. We have had to rewrite these entries with only the last component of the current triple in our metadata extractor.

3) Files in the parents block must be specified by their base names only. Currently, they are being specified using the full path to the temporary grid location. (you just addressed this one)

4) Finally, we currently use a custom module to fill metadata - in particular, we encode information about fcl files for simulation in the parents block that is being overwritten by these changes. In general, we need some way to override any fields that FileCatalogMetadata automatically fills.

We weren't aware of these metadata changes though looking at the release notes there is mention of some of these changes which appear to be addressing some outstanding bugs? We have a workaround for some of these steps but if there are any that appear to be "quick" fixes for the next release that'd be great - but we also don't want to hold up the release much since we're anticipating the bug fix for our calibration that addresses some memory issues we were having therein.

Thanks for the quick fix to the "parents" metadata issue. That was a real headache.

#3 - 08/20/2014 10:34 AM - Christopher Green

- Status changed from Resolved to Feedback

1. I was given to understand that unrecognized metadata items were ignored: is this not the case?
2. As with the above, and other information not provided by SAM: it is left to the experiment to decide exactly what information gets propagated to SAM and in what form. In other words, it was expected (at least by us), that some manipulation of the data would be required before shipping to SAM.
3. Fixed, as you say.
4. For most items, the metadata we fill are overridable by later entries. For automatically-generated stream-specific information, however (such as first_event, parents, etc.), they are done last and therefore cannot currently be overridden. In order to override the parents you'll have to use the parameters mentioned above.

However, after some overnight thought, I think it would be a better solution to change the system such that stream-specific user metadata (produced by a FileCatalogMetadataPlugin) automatically override items produced by the module, and then those parameters would be unnecessary. Does that seem like a reasonable thing to do?

#4 - 09/03/2014 09:35 AM - Christopher Green

- Status changed from Feedback to Resolved

Resolved with [cdf0d567d0df609791c7099ab0da335431e4cd8b](#): plugin-originated metadata items override module-originated items.

#5 - 09/03/2014 09:35 AM - Christopher Green

- Status changed from Resolved to Closed