

art - Support #4895

No way to filter runs/subruns out of output file

11/01/2013 05:45 PM - Christopher Backhouse

Status:	Closed	Start date:	11/01/2013
Priority:	Normal	Due date:	
Assignee:	Christopher Green	% Done:	100%
Category:	Event Loop	Estimated time:	0.00 hour
Target version:		Spent time:	1.00 hour
Scope:	Internal	SSI Package:	art
Experiment:	NOvA		
Description			
EDFilters can return false from their filter() function to cause downstream modules to not process this event. With the addition of a "SelectEvents" line in the RootOutput configuration one can also prevent those events going to disk.			
begin/end Run/Subrun also have boolean return values. Presumably these also prevent downstream modules from processing the entire [sub]run. (Actually, how does this work for endRun, where it's too late?). But I don't see any way to also filter those runs from disk. I tested that SelectEvents doesn't do so.			
Is it possible to remove these runs from the output file? If so, how? If not, can SelectEvents be expanded to do so, or SelectRun and SelectSubrun added?			
Thanks - Chris			
Related issues:			
Related to art - Feature #4916: Implement SelectRuns and SelectSubRuns		Feedback	11/06/2013
Related to art - Support #10545: Returning false from EDFilter::beginSubRun()...		Closed	10/15/2015

History

#1 - 11/04/2013 12:12 PM - Christopher Green

- Tracker changed from Bug to Support
- Category set to Event Loop
- Status changed from New to Resolved
- Assignee set to Christopher Green
- % Done changed from 0 to 100
- Experiment NOvA added
- Experiment deleted (-)
- SSI Package art added
- SSI Package deleted ()

We believe that, for the default file mode of, "FULLMERGE," setting services.scheduler.handleEmptySubRuns and services.scheduler.handleEmptyRuns both to false should give you the behavior you desire. If this is not the case, please submit a bug report for same.

#2 - 11/04/2013 02:00 PM - Christopher Backhouse

These don't appear to work.
The names seem confusing: the runs or subruns aren't necessarily empty, they've just been rejected by a filter.

What if I have "MERGE" mode (I didn't for my test)? (I'm not sure if it's necessary or not, but I inherited it from Issue [#4492](#)).

#3 - 11/05/2013 12:05 PM - Christopher Green

Unfortunately, you are perfectly correct: the decision is input-based, not output based. My apologies for that oversight in analysis.

That leaves you with the following: *reading* the file so written will avoid going through the run / subrun cycle for those without events, but you're still stuck putting them into the output file. You'd need to read it in FULLMERGE mode though, with the previously-mentioned parameters set -- MERGE won't do anything.

On a more general note, the file modes were inherited from CMS, and we are overdue for collecting requirements on the kind of things people want to see for what they need, rather than rummaging through the jumble of someone else's requirements.

#4 - 11/05/2013 01:44 PM - Christopher Backhouse

OK, those do indeed prevent the beginRun etc handlers being called.

Unfortunately, my file with a large number of filtered-out runs still takes a very long time to start processing the non-filtered run. So some non-trivial work is still being done with the filtered runs.

I broke in with gdb, and got one trace that looked like this (loading TTrees out of the file)

```
#25 0x00002aaaab7fcdd9 in TTree::GetEntry (this=0x2aab00b31710, entry=0,
getall=0)
at /archives/fwdev/products/root/v5_34_09/source/root/tree/tree/src/TTree.cxx:4918
#26 0x00002aaaafbe2ffe in art::input::getEntry (tree=0x2aab00b31710,
entryNumber=0)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/IO/Root/RootTree.cc:158
#27 0x00002aaaafbb8d3f in art::RootInputFile::RootInputFile (
this=0x2aab00988ca0, fileName=..., catalogName=...,
processConfiguration=..., logicalFileName=..., filePtr=...,
origEventID=..., eventsToSkip=0, whichSubRunsToSkip=..., fcip=...,
treeCacheSize=0, treeMaxVirtualSize=-1,
processingMode=art::InputSource::RunsSubRunsAndEvents, forcedRunOffset=0,
whichEventsToProcess=..., noEventSort=false, groupSelectorRules=...,
---Type <return> to continue, or q <return> to quit---
dropMergeable=false, duplicateChecker=..., dropDescendants=true)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/IO/Root/RootInputFile.cc:147
#28 0x00002aaaafbcceeda in art::RootInputFileSequence::initFile (
this=0x2aab008aec20, skipBadFiles=false)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/IO/Root/RootInputFileSequence.cc:238
#29 0x00002aaaafbcdd7e in art::RootInputFileSequence::RootInputFileSequence (
this=0x2aab008aec20, pset=..., catalog=..., fcip=...,
pMode=art::InputSource::RunsSubRunsAndEvents, mpr=..., processConfig=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/IO/Root/RootInputFileSequence.cc:90
#30 0x00002aab04009eff in art::RootInput::RootInput (this=0x2aab008ae440,
pset=..., desc=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/IO/Root/RootInput_source.cc:37
#31 0x00002aab0400b0da in make (ps=..., desc=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/IO/Root/RootInput_source.cc:250
#32 0x00002aab19cbb58 in art::InputSourceFactory::make (conf=..., desc=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/Core/InputSourceFactory.cc:52
#33 0x00002aaaab0c8081b in (anonymous namespace)::makeInput (params=...,
processName=..., preg=..., areg=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/EventProcessor/EventProcessor.cc:103
#34 0x00002aaaab0c81c53 in art::EventProcessor::EventProcessor (this=0x1a8bd10,
pset=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/EventProcessor/EventProcessor.cc:232
#35 0x00002aaaaab487de in art::run_art_common_ (main_pset=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/Art/run_art.cc:290
#36 0x00002aaaaab4790a in art::run_art(int, char**, boost::program_options::options_description&,
std::vector<std::unique_ptr<art::OptionsHandler, std::default_delete<art::OptionsHandler> >,
std::allocator<std::unique_ptr<art::OptionsHandler,
std::default_delete<art::OptionsHandler> > >>&&) (
argc=4, argv=0x7ffffff7078, in_desc=..., lookupPolicy=...,
handlers=<unknown type in /nusoft/app/externals/art/v1_08_04/slf5.x86_64.nu.e4.debug/lib/libart_Framework_Art.so, CU 0x124ec1, DIE 0x19acfb>)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/Art/run_art.cc:173
#37 0x00002aaaaab44ab4 in artapp (argc=4, argv=0x7ffffff7078)
at /archives/fwdev/products/art_suite/v1_08_04/build-Linux64bit+2.6-2.5/art-nu-e4-debug/art/Framework/Art/artapp.cc:47
#38 0x0000000000401af3 in main (argc=4, argv=0x7ffffff7078)
at /archives/fwdev/products/art_suite/v1_08_04/build-Linux64bit+2.6-2.5/art-nu-e4-debug/art/Framework/Art/nova.cc:8
```

and plenty that looked like this (parsing fcls?).

```
fhicl::parse_document (s=..., result=...)
---Type <return> to continue, or q <return> to quit---
at /archives/fwdev/products/art_suite/v1_08_04/include/fhiclcpp/v2_17_10/include/fhiclcpp/parse.h:43
#200 0x00002aaaab42be36d in fhicl::make_ParameterSet (str=..., ps=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/fhicl-cpp/fhiclcpp/make_ParameterSet.cc:109
#201 0x00002aaaafbb91a6 in art::RootInputFile::RootInputFile (
this=0x2aab00988ca0, fileName=..., catalogName=...,
processConfiguration=..., logicalFileName=..., filePtr=...,
origEventID=..., eventsToSkip=0, whichSubRunsToSkip=..., fcip=...,
treeCacheSize=0, treeMaxVirtualSize=-1,
processingMode=art::InputSource::RunsSubRunsAndEvents, forcedRunOffset=0,
whichEventsToProcess=..., noEventSort=false, groupSelectorRules=...,
```

dropMergeable=false, duplicateChecker=..., dropDescendants=true)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/IO/Root/RootInputFile.cc:192

#202 0x00002aaaafbcceeda in art::RootInputFileSequence::initFile (
this=0x2aab008aec20, skipBadFiles=false)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/IO/Root/RootInputFileSequence.cc:238

#203 0x00002aaaafbcdd7e in art::RootInputFileSequence::RootInputFileSequence (
this=0x2aab008aec20, pset=..., catalog=..., fcip=...,
pMode=art::InputSource::RunsSubRunsAndEvents, mpr=..., processConfig=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/IO/Root/RootInputFileSequence.cc:90

#204 0x00002aab04009eff in art::RootInput::RootInput (this=0x2aab008ae440,
pset=..., desc=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/IO/Root/RootInput_source.cc:37

#205 0x00002aab0400b0da in make (ps=..., desc=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/IO/Root/RootInput_source.cc:250

#206 0x00002aaab19cbb58 in art::InputSourceFactory::make (conf=..., desc=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/Core/InputSourceFactory.cc:52

#207 0x00002aaab0c8081b in (anonymous namespace)::makeInput (params=...,
processName=..., preg=..., areg=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/EventProcessor/EventProcessor.cc:103

#208 0x00002aaab0c81c53 in art::EventProcessor::EventProcessor (
this=0x1a8bd10, pset=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/EventProcessor/EventProcessor.cc:232

#209 0x00002aaaab487de in art::run_art_common_ (main_pset=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/Art/run_art.cc:290
#210 0x00002aaaab4790a in art::run_art(int, char**, boost::program_options::options_description&, cet::filepath_maker&,
std::vector<std::unique_ptr<art::OptionsHandler, std::default_delete<art::OptionsHandler> >, std::allocator<std::unique_ptr<art::OptionsHandler,
std::default_delete<art::OptionsHandler> > >>&&) (
argc=4, argv=0x7ffffff7078, in_desc=..., lookupPolicy=...,
handlers=<unknown type in /nusoft/app/externals/art/v1_08_04/slf5.x86_64.nu.e4.debug/lib/libart_Framework_Art.so, CU 0x124ec1, DIE 0x19acfb>)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/Art/run_art.cc:173

#211 0x00002aaaab44ab4 in artapp (argc=4, argv=0x7ffffff7078)
at /archives/fwdev/products/art_suite/v1_08_04/build-Linux64bit+2.6-2.5/art-nu-e4-debug/art/Framework/Art/artapp.cc:47

#212 0x000000000401af3 in main (argc=4, argv=0x7ffffff7078)
at /archives/fwdev/products/art_suite/v1_08_04/build-Linux64bit+2.6-2.5/art-nu-e4-debug/art/Framework/Art/nova.cc:8

and this (building parameter sets?)

#9 0x00002aaab42be380 in fhicl::make_ParameterSet (str=..., ps=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/fhicl-cpp/fhiclcpp/make_ParameterSet.cc:109

#10 0x00002aaaafbb91a6 in art::RootInputFile::RootInputFile (
this=0x2aab00988ca0, fileName=..., catalogName=...,
processConfiguration=..., logicalFileName=..., filePtr=...,
origEventID=..., eventsToSkip=0, whichSubRunsToSkip=..., fcip=...,
treeCacheSize=0, treeMaxVirtualSize=-1,
processingMode=art::InputSource::RunsSubRunsAndEvents, forcedRunOffset=0,
whichEventsToProcess=..., noEventSort=false, groupSelectorRules=...,
dropMergeable=false, duplicateChecker=..., dropDescendants=true)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/IO/R---Type <return> to continue, or q <return> to quit---
oot/RootInputFile.cc:192

#11 0x00002aaaafbcceeda in art::RootInputFileSequence::initFile (
this=0x2aab008aec20, skipBadFiles=false)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/IO/Root/RootInputFileSequence.cc:238

#12 0x00002aaaafbcdd7e in art::RootInputFileSequence::RootInputFileSequence (
this=0x2aab008aec20, pset=..., catalog=..., fcip=...,
pMode=art::InputSource::RunsSubRunsAndEvents, mpr=..., processConfig=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/IO/Root/RootInputFileSequence.cc:90

#13 0x00002aab04009eff in art::RootInput::RootInput (this=0x2aab008ae440,
pset=..., desc=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/IO/Root/RootInput_source.cc:37

#14 0x00002aab0400b0da in make (ps=..., desc=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/IO/Root/RootInput_source.cc:250

#15 0x00002aaab19cbb58 in art::InputSourceFactory::make (conf=..., desc=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/Core/InputSourceFactory.cc:52

#16 0x00002aaab0c8081b in (anonymous namespace)::makeInput (params=...,
processName=..., preg=..., areg=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/EventProcessor/EventProcessor.cc:103

#17 0x00002aaab0c81c53 in art::EventProcessor::EventProcessor (this=0x1a8bd10,
pset=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/EventProcessor/EventProcessor.cc:232

#18 0x00002aaaab487de in art::run_art_common_ (main_pset=...)
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/Art/run_art.cc:290

```
#19 0x00002aaaaab4790a in art::run_art(int, char**, boost::program_options::options_description&, cet::filepath_maker&,
std::vector<std::unique_ptr<art::OptionsHandler, std::default_delete<art::OptionsHandler> >, std::allocator<std::unique_ptr<art::OptionsHandler,
std::default_delete<art::OptionsHandler> > >>&&) (
argc=4, argv=0x7ffffff7078, in_desc=..., lookupPolicy=...,
handlers=<unknown type in /nusoft/app/externals/art/v1_08_04/slf5.x86_64.nu.e4.debug/lib/libart_Framework_Art.so, CU 0x124ec1, DIE 0x19acfb>
at /archives/fwdev/products/art_suite/v1_08_04/source/art/art/Framework/Art/run_art.cc:173
#20 0x00002aaaaab44ab4 in artapp (argc=4, argv=0x7ffffff7078)
at /archives/fwdev/products/art_suite/v1_08_04/build-Linux64bit+2.6-2.5/art-nu-e4-debug/art/Framework/Art/artapp.cc:47
#21 0x000000000401af3 in main (argc=4, argv=0x7ffffff7078)
at /archives/fwdev/products/art_suite/v1_08_04/build-Linux64bit+2.6-2.5/art-nu-e4-debug/art/Framework/Art/nova.cc:8
```

Is this so that the final file can have the configurations of all the input files included in it? Any optimizations available here?

#5 - 11/05/2013 01:45 PM - Christopher Backhouse

If there is a revamp of the file modes, I would suggest SelectRuns and SelectSubRuns to mirror SelectEvents.

#6 - 11/06/2013 09:50 AM - Christopher Green

I can't honestly think of an optimization here, sorry. I would suggest though, that you put in the SelectRuns, *etc.* suggestion as a separate feature request so it doesn't get lost.

#7 - 03/03/2014 11:45 AM - Lynn Garren

- *Status changed from Resolved to Closed*

#8 - 11/02/2015 11:51 AM - Christopher Green

- *Related to Support #10545: Returning false from EDFilter::beginSubRun() does not skip the events added*