

artdaq - Support #21626

Notes on where data logging (art) processes are spending their time

01/04/2019 05:25 PM - Kurt Biery

Status:	Assigned	Start date:	01/04/2019
Priority:	Normal	Due date:	
Assignee:	Kurt Biery	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Co-Assignees:	
Experiment:	-		
Description			
<p>In some of my recent artdaq-demo tests on mu2eddaq01, I noticed that the event rate through the system would sometimes drop to zero for several seconds. After some poking around, I realized/remembered that this happens when one output file is closed and another is opened. Of course, it would be interesting to know why the closing of the files takes as long as it does.</p> <p>This may be somewhat related to understanding where data logging processes spend their time in general. At protoDUNE, there was a sense among some of our colleagues that the performance of disk-writing artdaq/art processes could have been incrementally better.</p>			

History

#1 - 01/04/2019 05:33 PM - Kurt Biery

I'm using the code on the feature/TracesForTimeSpentInArt branches in the artdaq_core and artdaq repositories.

With this code, I've enabled the following TRACE levels:

- SharedMemoryManager 4
- SharedMemoryReader 4
- RootDAQOutFile 4

I'm sure that there will be other TRACE levels that will be useful.

#2 - 01/04/2019 06:01 PM - Kurt Biery

A simple reminder is that data logging processes that are asked to send data to a Dispatcher process will incur the cost of doing that. In an artdaq-demo run on mu2eddaq01 with the mediumsystem_with_routing_master configuration, in which the event size is approximately 9 MB, the event rate is 10 Hz, and there are 4 EventBuilders, the following TRACE messages were observed:

```
5990 01-04 18:00:37.165020 30590 30590 5 SharedMemoryManager d04 . Before memcpy in Read(), size is 9006112
6010 01-04 18:00:37.163220 30590 30590 5 RootNetOutput d12 . End: RootNetOutput::write(const EventPrincipal& ep)
6011 01-04 18:00:37.163209 30590 30590 5 RootNetOutput d12 . RootNetOutput::write(const EventPrincipal& ep): Message sent.
6012 01-04 18:00:37.162422 30590 30590 5 DataLogger1_art1_DataSenderManager d05 . sendFragment: Done sending fragment 349 to dest=-1111
6013 01-04 18:00:37.162374 30590 30590 5 DataLogger1_art1_DataSenderManager d05 . sendFragment: sending metrics
6023 01-04 18:00:37.161552 30590 30590 5 DataLogger1_art1_TCPSocketTransfer d12 . transfer_between_14_and_15_SEND: sendFragment returning Success
6036 01-04 18:00:37.157580 30590 30590 5 DataLogger1_art1_TCPSocketTransfer d12 . transfer_between_14_and_15_SEND: sendFragment begin send of fragment with sequenceID=349
6067 01-04 18:00:37.149189 30590 30590 5 DataLogger1_art1_DataSenderManager d04 . sendFragment: Sending fragment with seqId 349 to destination 15 (broadcast)
6068 01-04 18:00:37.143089 30590 30590 5 NetMonTransportService dbg . Sending message with sequenceID=349, type=226, length=9006336
6069 01-04 18:00:37.143058 30590 30590 5 RootNetOutput d12 . RootNetOutput::write(const EventPrincipal& ep): Sending a message ...
6258 01-04 18:00:37.121672 30590 30590 5 RootNetOutput d12 . RootNetOutput::write(const EventPrincipal& ep): Finished streaming History.
6259 01-04 18:00:37.121667 30590 30590 5 RootNetOutput d12 . RootNetOutput::write(const EventPrincipal& ep): Streaming History ...
6260 01-04 18:00:37.121666 30590 30590 5 RootNetOutput d12 . RootNetOutput::write(const EventPrincipal& ep): Finished streaming EventAuxiliary.
6261 01-04 18:00:37.121661 30590 30590 5 RootNetOutput d12 . RootNetOutput::write(const EventPrincipal& ep): Streaming EventAuxiliary ...
```

```

6262 01-04 18:00:37.121660 30590 30590 5 RootNetOutput d12 . RootNetOutput::write(const EventPrincipal& ep): Finished streaming SubRunAuxiliary.
6263 01-04 18:00:37.121654 30590 30590 5 RootNetOutput d12 . RootNetOutput::write(const EventPrincipal& ep): Streaming SubRunAuxiliary ...
6264 01-04 18:00:37.121653 30590 30590 5 RootNetOutput d12 . RootNetOutput::write(const EventPrincipal& ep): Finished streaming RunAuxiliary.
6265 01-04 18:00:37.121637 30590 30590 5 RootNetOutput d12 . RootNetOutput::write(const EventPrincipal& ep): Streaming RunAuxiliary ...
6266 01-04 18:00:37.121636 30590 30590 5 RootNetOutput d12 . RootNetOutput::write(const EventPrincipal& ep): Finished streaming message type code.
6267 01-04 18:00:37.121634 30590 30590 5 RootNetOutput d12 . RootNetOutput::write(const EventPrincipal& ep): Streaming message type code ...
6268 01-04 18:00:37.121621 30590 30590 5 RootNetOutput d12 . Begin: RootNetOutput::write(const EventPrincipal& ep)
6269 01-04 18:00:37.121577 30590 30590 5 DataLogger1_art1_RootDAQOutFile d04 . End of RootDAQOutFile::writeOne
6270 01-04 18:00:37.121522 30590 30590 5 DataLogger1_art1_RootDAQOutFile d04 . End of RootDAQOutFile::fillBranches
6271 01-04 18:00:37.121516 30590 30590 5 DataLogger1_art1_RootDAQOutFile d04 . RootDAQOutFile::fillBranches after fillTree call
6272 01-04 18:00:37.110354 30590 30590 5 DataLogger1_art1_RootDAQOutFile d04 . RootDAQOutFile::fillBranches before fillTree call
6273 01-04 18:00:37.110302 30590 30590 5 DataLogger1_art1_RootDAQOutFile d04 . Start of RootDAQOutFile::fillBranches
6274 01-04 18:00:37.110290 30590 30590 5 DataLogger1_art1_RootDAQOutFile d04 . Start of RootDAQOutFile::writeOne
6431 01-04 18:00:37.094885 30590 30590 5 SharedMemoryManager d04 . After memcpy in Read()
6546 01-04 18:00:37.087599 30590 30590 5 SharedMemoryManager d04 . Before memcpy in Read(), size is 9006112

```

The additional TRACE levels that were useful here were:

- RootNetOutput level 12
- DataLogger1_art1_DataSenderManager levels 4 and 5
- DataLogger1_art1_TCPsocketTransfer level 12

#3 - 01/04/2019 06:08 PM - Kurt Biery

Some general notes on tests with the current mediumsystm_with_routing_master configuration on mu2edaq01...

In longer runs with the sending of events to the Dispatcher turned ON, the DataLogger art subprocess uses ~85% of a CPU before the first file close and ~100% thereafter (when writing data; when the file close happens, the CPU usage is low). What I believe is happening is that the data logging gets behind when the first file close happens and then never recovers because the data logging process is running close to capacity.

With the sending of events turned OFF, the data logging process takes ~35% of a CPU before the first file close, and returns to that CPU usage after it clears out the backlog after each file close.

#4 - 01/04/2019 06:26 PM - Kurt Biery

Here are some TRACE messages related to file closing:

```

[biery@mu2edaq01 data]$ tshow | grep 21225 | grep -v touchBuffer_ | grep -v checkBuffer_ | egrep 'writeTTree|output file' | tdelta -ct 1
594 01-04 18:22:19.189561 21225 21225 13 DataLogger1_art1_RootDAQOut nfo . finishEndFile: Closed output file "/scratch/biery/data/artdaqdemo_r000236_sr01_4_dll.root"
598 01-04 18:22:19.094594 21225 21225 13 DataLogger1_art1_RootDAQOutFile d04 . End of RootDAQOutFile::writeTTrees
599 01-04 18:22:19.090958 21225 21225 13 DataLogger1_art1_RootDAQOutFile d04 . RootDAQOutFile::writeTTrees after writing parentageTree_
600 01-04 18:22:19.057903 21225 21225 3 DataLogger1_art1_RootDAQOutFile d04 . RootDAQOutFile::writeTTrees after writing fileIndexTree_
601 01-04 18:22:19.057780 21225 21225 3 DataLogger1_art1_RootDAQOutFile d04 . RootDAQOutFile::writeTTrees after writing metaDataTree_
850 01-04 18:22:04.522872 21225 21225 2 DataLogger1_art1_RootDAQOutFile d04 . Start of RootDAQOutFile::writeTTrees
26393 01-04 18:21:43.540975 21225 21225 3 DataLogger1_art1_RootDAQOut nfo . doOpenFile: Opened output file with pattern "/scratch/biery/data/artdaqdemo_r%06r_sr%02s_%#_dll.root"
26395 01-04 18:21:43.515176 21225 21225 2 DataLogger1_art1_RootDAQOut nfo . finishEndFile: Closed output file "/scratch/biery/data/artdaqdemo_r000236_sr01_3_dll.root"
26430 01-04 18:21:43.410812 21225 21225 3 DataLogger1_art1_RootDAQOutFile d04 . End of RootDAQOutFile::writeTTrees
26431 01-04 18:21:43.407629 21225 21225 3 DataLogger1_art1_RootDAQOutFile d04 . RootDAQOutFile::writeTTrees after writing parentageTree_
26463 01-04 18:21:43.374537 21225 21225 2 DataLogger1_art1_RootDAQOutFile d04 . RootDAQOutFile::wr

```

```

iteTTrees after writing fileIndexTree_
 26464 01-04 18:21:43.374440 21225 21225 2 DataLogger1_art1_RootDAQOutFile d04 . RootDAQOutFile::wr
iteTTrees after writing metaDataTree_
 69696 01-04 18:21:16.863987 21225 21225 14 DataLogger1_art1_RootDAQOutFile d04 . Start of RootDAQOu
tFile::writeTTrees
 394614 01-04 18:19:59.651397 21225 21225 14 DataLogger1_art1_RootDAQOut nfo . doOpenFile: Opened
output file with pattern "/scratch/biery/data/artdaqdemo_r%06r_sr%02s_%#_dll.root"
 394617 01-04 18:19:59.625461 21225 21225 2 DataLogger1_art1_RootDAQOut nfo . finishEndFile: Clo
sed output file "/scratch/biery/data/artdaqdemo_r000236_sr01_2_dll.root"
 394626 01-04 18:19:59.539136 21225 21225 3 DataLogger1_art1_RootDAQOutFile d04 . End of RootDAQOutF
ile::writeTTrees
 394627 01-04 18:19:59.534528 21225 21225 3 DataLogger1_art1_RootDAQOutFile d04 . RootDAQOutFile::wr
iteTTrees after writing parentageTree_
 394629 01-04 18:19:59.509957 21225 21225 2 DataLogger1_art1_RootDAQOutFile d04 . RootDAQOutFile::wr
iteTTrees after writing fileIndexTree_
 394630 01-04 18:19:59.509856 21225 21225 2 DataLogger1_art1_RootDAQOutFile d04 . RootDAQOutFile::wr
iteTTrees after writing metaDataTree_
 456133 01-04 18:19:34.670945 21225 21225 0 DataLogger1_art1_RootDAQOutFile d04 . Start of RootDAQOu
tFile::writeTTrees
 781408 01-04 18:18:14.703257 21225 21225 3 DataLogger1_art1_RootDAQOut nfo . doOpenFile: Opened
output file with pattern "/scratch/biery/data/artdaqdemo_r%06r_sr%02s_%#_dll.root"
 781410 01-04 18:18:14.669476 21225 21225 2 DataLogger1_art1_RootDAQOut nfo . finishEndFile: Clo
sed output file "/scratch/biery/data/artdaqdemo_r000236_sr01_1_dll.root"
 781523 01-04 18:18:14.582063 21225 21225 3 DataLogger1_art1_RootDAQOutFile d04 . End of RootDAQOutF
ile::writeTTrees
 781524 01-04 18:18:14.578237 21225 21225 3 DataLogger1_art1_RootDAQOutFile d04 . RootDAQOutFile::wr
iteTTrees after writing parentageTree_
 781531 01-04 18:18:14.545453 21225 21225 2 DataLogger1_art1_RootDAQOutFile d04 . RootDAQOutFile::wr
iteTTrees after writing fileIndexTree_
 781532 01-04 18:18:14.545343 21225 21225 2 DataLogger1_art1_RootDAQOutFile d04 . RootDAQOutFile::wr
iteTTrees after writing metaDataTree_
 822268 01-04 18:18:01.162565 21225 21225 15 DataLogger1_art1_RootDAQOutFile d04 . Start of RootDAQOu
tFile::writeTTrees
1183490 01-04 18:16:29.221611 21225 21225 3 DataLogger1_art1_RootDAQOut nfo . doOpenFile: Opened
output file with pattern "/scratch/biery/data/artdaqdemo_r%06r_sr%02s_%#_dll.root"

```