

artdaq - Support #13657

Investigation: max disk-writing rate with RAM disk

08/24/2016 10:00 AM - Kurt Biery

Status:	Assigned	Start date:	08/24/2016
Priority:	Normal	Due date:	
Assignee:	Kurt Biery	% Done:	0%
Category:	Known Issues	Estimated time:	0.00 hour
Target version:		Co-Assignees:	
Experiment:	-		
Description			
As part of preparing for protoDUNE, we were asked to look into the maximum data rate that can be supported by a disk-writing Aggregator. Ron and Eric used the Mu2e DAQ cluster to run this test. The purpose of this Issue is simply to capture the notes/emails on this topic.			

History

#1 - 08/24/2016 10:01 AM - Kurt Biery

Yes, I was just thinking about that and looking into it...

New table:

no disk writing: 500 MB/s

Root file results, no compress:

ram disk writing:440 MB/s

hdd disk writing 77 MB/s

Root file results, Compression:7 :

ram disk writing: 40 MB/s

hdd disk writing: 40 MB/s

Binary file results:

ram disk writing:200 MB/s

hdd disk writing: 86 MB/s (probably didn't wait long enough for a (lower) steady-state rate

- 4x40x2 system
- why is binary slower than root??
- no periodic drop_caches

Kurt Biery wrote on 08/23/2016 01:33 PM:

Or, maybe we have ROOT compression turned on. I don't remember what the default ROOT compression level is, but in some cases, we have explicitly turned it off. ("compressionLevel: 0" in the RootOutput configuration <<https://cdcv.s.fnal.gov/redmine/projects/artdaq-demo/repository/revisions/develop/entry/tools/generateAggregatorMain.rb#L36>>).
Kurt

On 8/23/16 1:29 PM, Ron Rechenmacher wrote:

no disk writing: 500 MB/s

Initial root file results:

ram disk writing: 40 MB/s

hdd disk writing: 40 MB/s

Initial binary file results:

ram disk writing:200 MB/s

hdd disk writing: 86 MB/s

Writing root files... I think we might somehow have put a spin-wait into it???? ;}

I'll try to look into it.

--Ron

Ron Rechenmacher wrote on 08/23/2016 12:41 PM:

Should have results shortly.

Kurt Biery wrote on 08/22/2016 01:12 PM:

Hi Ron, Eric, all,

protoDUNE folks are interested in how quickly we can write data to disk in /artdaq/. This is a very reasonable question. A suggestion was made to use a RAM disk and run some tests. We could try doing that on the DUNE 35-ton DAQ cluster, but there could be complications in getting things set up. How about the Mu2e Pilot system? I seem to recall Eric saying that the max rates that he mentions are based on tests with disk writing turned off. Could we create a RAM

disk on the relevant computer, turn disk-writing on, and see what rate can be supported by the mu2eddaq system with disk writing?
Thanks,
Kurt

--

Kurt Biery

/Applications Development and Systems Analysis Manager; Department Head/

Real-Time Systems Engineering Department
Fermi National Accelerator Laboratory
630 840 6663 office
www.fnal.gov <<http://www.fnal.gov>>
biery@fnal.gov <<mailto:biery@fnal.gov>>

--

Ron Rechenmacher
Engineer
Fermi National Accelerator Laboratory
Batavia, IL 60510

#2 - 08/24/2016 10:14 AM - Kurt Biery

Sent to dune-proto-sp-daq@fnal.gov and artdaq-developers@fnal.gov:

Sorry, I should have also mentioned that the the event size in these tests, as delivered to the Data Logger, was ~38 MB. This is not an exact match to the 50 MB that was requested, but it is close.
Kurt

On 8/24/16 9:42 AM, Kurt Biery wrote:

Hi everyone,
As requested in the meeting on Monday, we've run some tests in which we focused on the performance of the disk writing in the Data Logging Aggregator (i.e. Data Logger). To do this, we used a cluster of computers that we have for developing the Mu2e DAQ system and created a temporary artificial disk in memory (i.e. RAM disk) to remove the effect of disk performance as much as possible.

Here are the results (the artdaq results are the data rates delivered to the Data Logger):

writes to the RAM disk outside of artdaq - 3 GB/s
maximum rate delivered to the Data Logger with disk writing turned off - 500 MB/s
max rate to disk with ROOT compression level of 7 (maximum) - 40 MB/s
max rate to disk with ROOT compression turned off (value of 0) - 440 MB/s

The maximum rate without disk writing is, of course, a function of the artdaq configuration that we used, and it provides the upper bound on what we can expect from the test with disk-writing turned on in this configuration.

Clearly, ROOT compression has a noticeable impact. In studies for LArIAT, John Freeman found that a ROOT compression setting of 3 provides the best balance between reducing the data size and not consuming too much CPU. Given the plan for compressing the data in the RCE in protoDUNE SP, I believe that we should focus on the result without ROOT compression here.

This test used an artdaq configuration in which there were 10 EventBuilder processes running on each of 4 computers (40 EBs total) delivering complete events to an Aggregator process running on a fifth computer. The EB computers were connected to a 10 Gb switch with 10 Gb NICs, and the Aggregator computer was connected to the switch with 2 10 Gbps links.

Thanks to Ron Rechenmacher and Eric Flumerfelt for running the tests.
Kurt

--

Kurt Biery
Applications Development and Systems Analysis Manager; Department Head

Real-Time Systems Engineering Department
Fermi National Accelerator Laboratory
630 840 6663 office
www.fnal.gov
biery@fnal.gov

#3 - 11/14/2016 01:19 PM - Eric Flumerfelt

- Target version set to 575

#4 - 01/13/2017 03:47 PM - Eric Flumerfelt

- *Category set to Known Issues*

- *Target version deleted (575)*