

Optimizing CMS Data Formats for Analysis

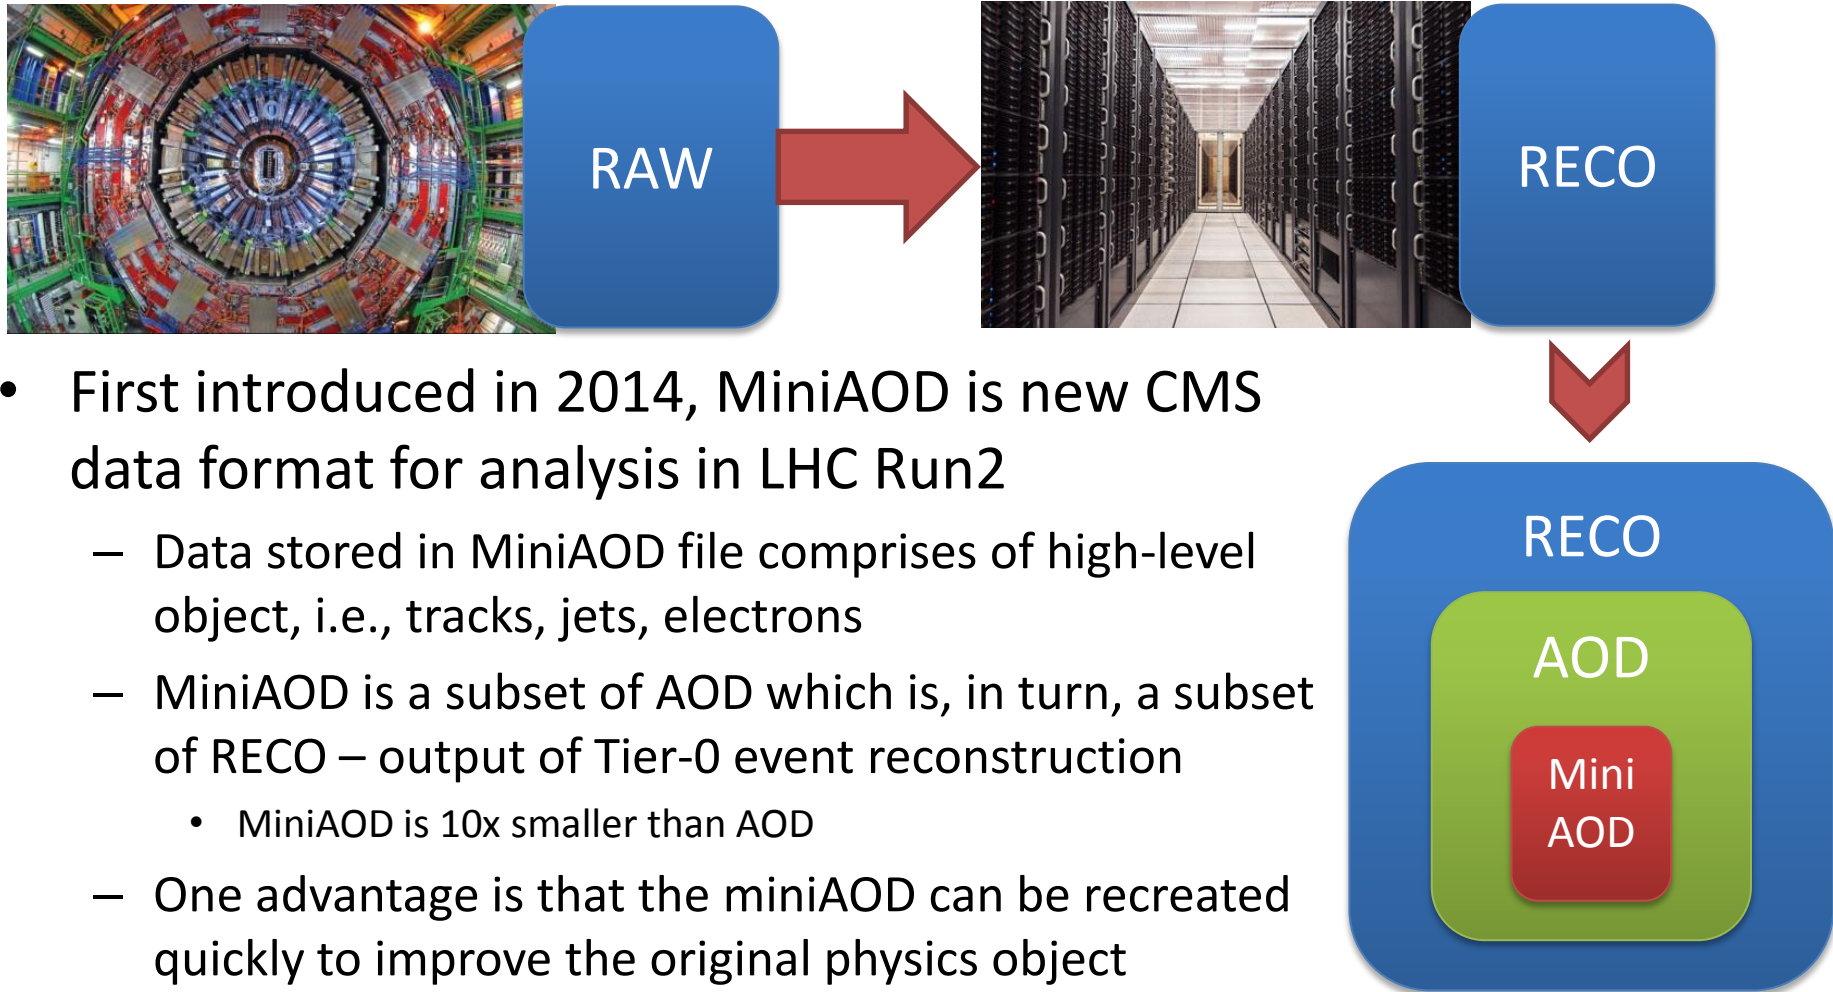
Peerut Boonchokchuay

August 11th, 2015

Outline

- Introduction to MiniAOD
- Objectives for Project
- Method Developed
- Example Results
- Conclusion and Next Steps

CMS MiniAOD Data Format

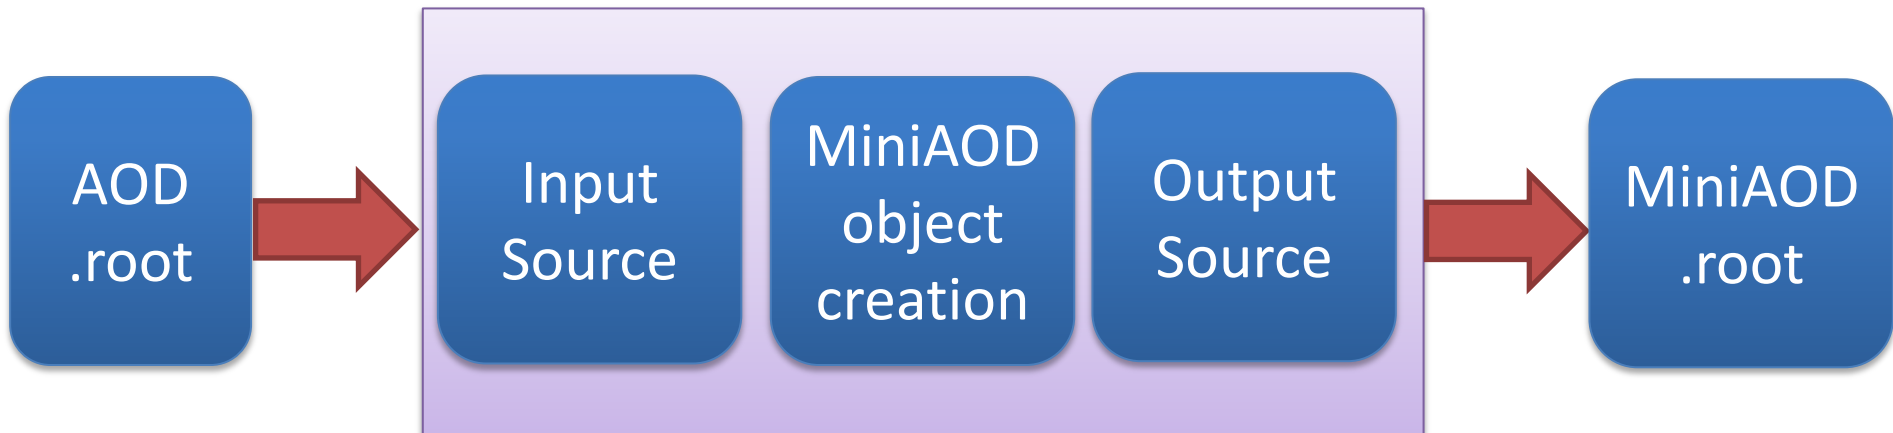


- First introduced in 2014, MiniAOD is new CMS data format for analysis in LHC Run2
 - Data stored in MiniAOD file comprises of high-level object, i.e., tracks, jets, electrons
 - MiniAOD is a subset of AOD which is, in turn, a subset of RECO – output of Tier-0 event reconstruction
 - MiniAOD is 10x smaller than AOD
 - One advantage is that the miniAOD can be recreated quickly to improve the original physics object definitions using the latest algorithms

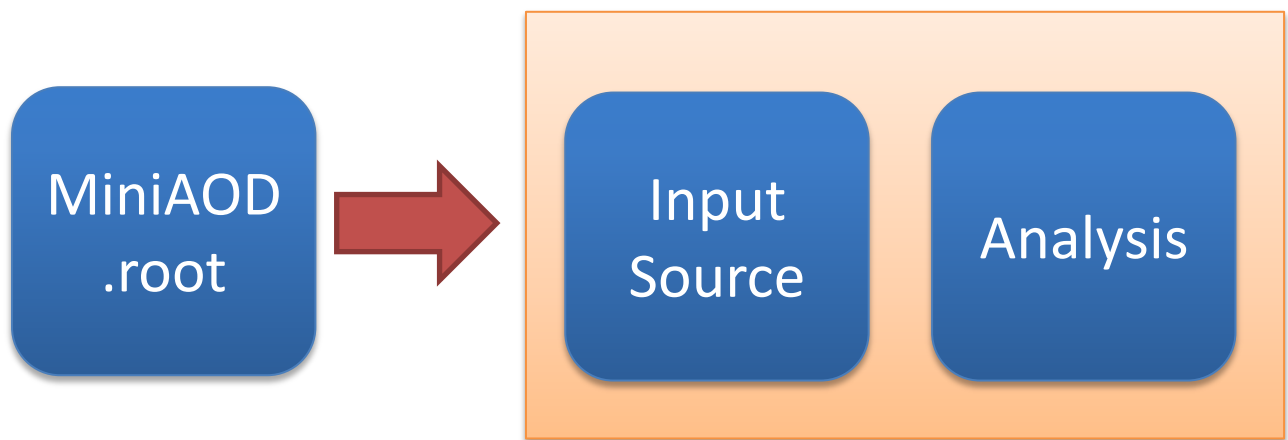
Objectives

- As with any analysis formats, users can benefit from faster processing times and smaller data sizes (to fit more data on your laptop!)
- To investigate ways to improve the MiniAOD performance in terms of:
 1. Time to read events from file
 2. Size of output filewhile withstanding:
 1. Job memory usage (RSS of CMSSW application)
 2. Time to write events to file

Step 1: Create MiniAOD



Step 2: Analyze MiniAOD



MiniAOD
modules
and
parameters
are based on
ROOT

Method

Study MiniAOD Creation Process

- Execution time
 - Write time
 - Compression time
- Output file size
- Memory usage

Study MiniAOD Reading Process

- Execution time
 - Read time
 - Decompression time

Adjust Parameter Settings

- Number of events
- Basket Size
- Compression Algorithm
- Compression Level
-

Measure Performance

- Tools
 - igProf
 - RSS monitoring
- In terms of:
 - CPU time
 - memory usage

Counter: PERF_TICKS, first 1000 entries

Sorted by cumulative cost

(Sort by self cost)

Rank	Total %	Cumulative	Symbol name
<u>1</u>	100.00	212.30	<u><spontaneous></u>
<u>3</u>	99.99	212.29	<u>__libc_start_main</u>
<u>2</u>	99.99	212.29	<u>__start</u>
<u>4</u>	99.33	210.89	<u>main</u>
<u>5</u>	99.28	210.78	<u>main::#1::operator()() const</u>
<u>6</u>	77.01	163.50	<u>edm::EventProcessor::runToCompletion()</u>
<u>10</u>	69.24	147.01	<u>edm::EventProcessor::readAndProcessEvent()</u>
<u>9</u>	69.24	147.01	<u>statemachine::HandleEvent::readAndProcessEvent()</u>
<u>8</u>	69.24	147.01	<u>statemachine::HandleEvent::HandleEvent(boost::statechart::state<statemachine::HandleEvent, statemachine::Ha</u>
<u>7</u>	69.24	147.01	<u>boost::statechart::state<statemachine::HandleEvent, statemachine::HandleLumis, boost::mpl::list<mpl::na, m</u>
<u>13</u>	69.24	147.00	<u>edm::EventProcessor::processEventsForStreamAsync(unsigned int, std::atomic<bool>*)</u>
<u>12</u>	69.24	147.00	<u>edm::StreamProcessingTask::execute()</u>
<u>11</u>	69.24	147.00	<u>tbb::internal::custom_scheduler<tbb::internal::IntelSchedulerTraits>::local_wait_for_all(tbb::task&, tbb::t</u>
<u>14</u>	69.02	146.54	<u>edm::EventProcessor::processEvent(unsigned int)</u>
<u>17</u>	68.47	145.37	<u>void edm::StreamSchedule::processOneEvent<edm::OccurrenceTraits<edm::EventPrincipal, (edm::BranchActionType</u>
<u>16</u>	68.47	145.37	<u>decltype ({parm#1}()) edm::convertException::wrap<void edm::StreamSchedule::processOneEvent<edm::Occurrence</u>
<u>15</u>	68.47	145.37	<u>void edm::StreamSchedule::processOneEvent<edm::OccurrenceTraits<edm::EventPrincipal, (edm::BranchActionType</u>
<u>19</u>	68.46	145.35	<u>decltype ({parm#1}()) edm::convertException::wrap<bool edm::Worker::doWork<edm::OccurrenceTraits<edm::Event</u>
<u>18</u>	68.46	145.35	<u>bool edm::Worker::doWork<edm::OccurrenceTraits<edm::EventPrincipal, (edm::BranchActionType)1> >(edm::Occurr</u>
<u>20</u>	68.45	145.33	<u>void edm::Path::processOneOccurrence<edm::OccurrenceTraits<edm::EventPrincipal, (edm::BranchActionType)1> ></u>
<u>21</u>	68.45	145.31	<u>decltype ({parm#1}()) edm::convertException::wrap<void edm::Path::processOneOccurrence<edm::OccurrenceTrait</u>

Example of igProf Results

```

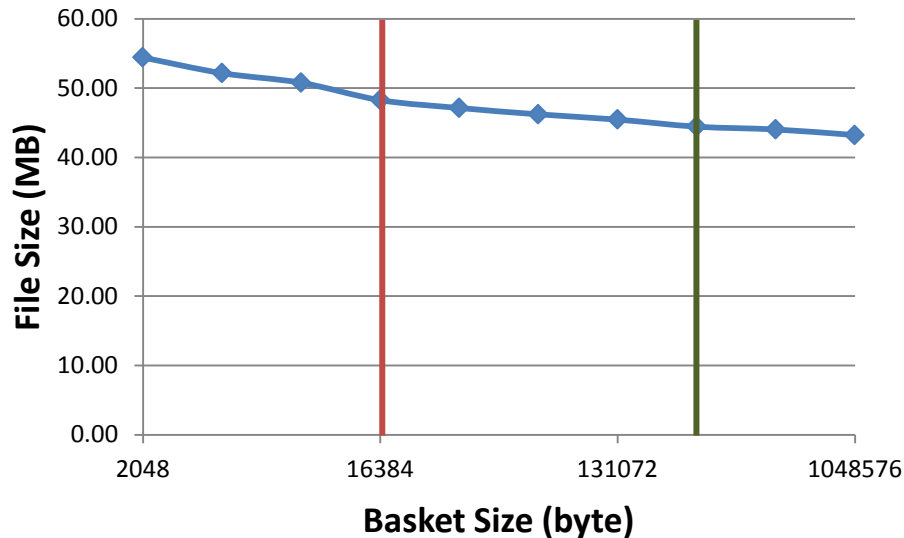
85 8.25 17.91 boost::statechart::simple state<statemachine::FirstLumi, statemachine::HandleLumis, boost::mpl::list<
86 7.57 16.43 TMVA::Node::ReadXML(void*, unsigned int)'4
87 6.96 15.10 edm::InputProductHolder::resolveProduct (edm::ProductHolderBase::ResolveStatus&, bool, edm::ModuleCal
88 6.95 15.08 edm::reftobase::IndirectVectorHolder<reco::Candidate>::at(unsigned long) const
89 6.92 15.02 edm::EventPrincipal::readFromSource (edm::ProductHolderBase const&, edm::ModuleCallingContext const*)
90 6.85 14.87 TMVA::Node::ReadXML(void*, unsigned int)'5
91 6.80 14.75 edm::DelayedReader::getProduct(edm::BranchKey const&, edm::EDProductGetter const*)
92 6.77 14.70 edm::RootDelayedReader::getProduct (edm::BranchKey const&, edm::EDProductGetter const*) const
93 6.77 14.69 @liblzma.so.5.0.3+86865}
94 6.46 14.03 void TMVA::Tools::ReadAttr<float>(void*, char const*, float&)
95 6.28 13.63 edm::RootDelayedReader::getProduct (edm::BranchKey const&, edm::EDProductGetter const*) const
96 6.24 13.63 edm::RootDelayedReader::getProduct (edm::BranchKey const&, edm::EDProductGetter const*) const
97 6.18 13.63 edm::RootDelayedReader::getProduct (edm::BranchKey const&, edm::EDProductGetter const*) const
98 5.99 13.63 edm::RootDelayedReader::getProduct (edm::BranchKey const&, edm::EDProductGetter const*) const
99 5.50 13.63 edm::RootDelayedReader::getProduct (edm::BranchKey const&, edm::EDProductGetter const*) const
100 5.45 13.63 edm::RootDelayedReader::getProduct (edm::BranchKey const&, edm::EDProductGetter const*) const
101 4.97 13.63 edm::RootDelayedReader::getProduct (edm::BranchKey const&, edm::EDProductGetter const*) const
102 4.88 13.63 edm::RootDelayedReader::getProduct (edm::BranchKey const&, edm::EDProductGetter const*) const
108 4.86 13.63 edm::RootDelayedReader::getProduct (edm::BranchKey const&, edm::EDProductGetter const*) const
107 4.86 13.63 edm::RootDelayedReader::getProduct (edm::BranchKey const&, edm::EDProductGetter const*) const
106 4.86 13.63 edm::RootDelayedReader::getProduct (edm::BranchKey const&, edm::EDProductGetter const*) const
105 4.86 10.54 boost::statechart::detail::reaction result boost::statechart::simple state<statemachine::handleFiles,
104 4.86 10.54 boost::statechart::simple state<statemachine::HandleRuns, statemachine::HandleFiles, statemachine::Ne
103 4.86 10.54 boost::statechart::simple state<statemachine::HandleLumis, statemachine::HandleRuns, statemachine::Fi

```


Rank	% total	Counts		Paths		Symbol name
		to / from this	Total	Including child / parent	Total	
	7.09	14.97	15.06	27	30	edm::InputProductHolder::resolveProduct (edm::Pr
[89]	7.09	0.05	14.93	27	27	edm::EventPrincipal::readFromSource_(edm::Produ
	6.98	14.73	14.74	25	26	edm::DelayedReader::getProduct(edm::BranchKey co
	0.04	0.08	0.09	3	4	edm::Principal::checkUniquenessAndType(edm::Wrap
	0.02	0.04	0.04	3	3	edm::BranchKey::BranchKey(edm::BranchDescription
	0.02	0.04	0.04	3	3	std:: Sp counted deleter<decltype(nullptr), edm
	0.01	0.02	0.02	2	2	edm::InputProductHolder::putProduct (std::unique
	0.01	0.01	5.54	1	494	init

Example Results: MiniAOD Reading Process

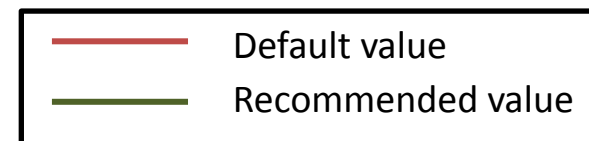
Output File Size vs Basket Size



Read MiniAod Time vs Basket Size




- By adjusting basket size from the default 16384 byte to 262144 byte, we can improve MiniAOD file size by 8% and read back time by 9%

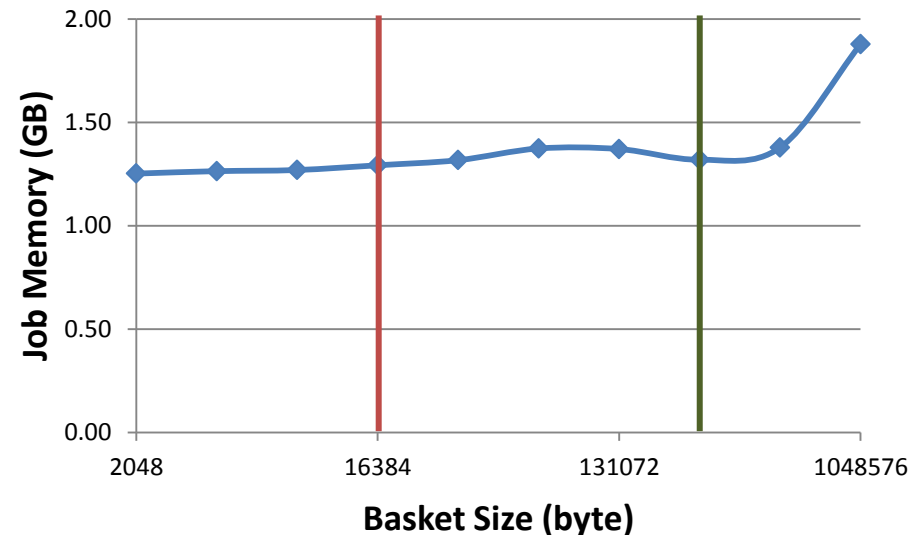


Example Results: MiniAOD Creation Process

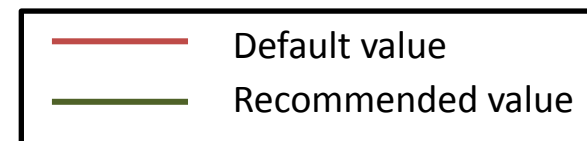
Output File Size vs Basket Size



Job Memory vs Basket Size

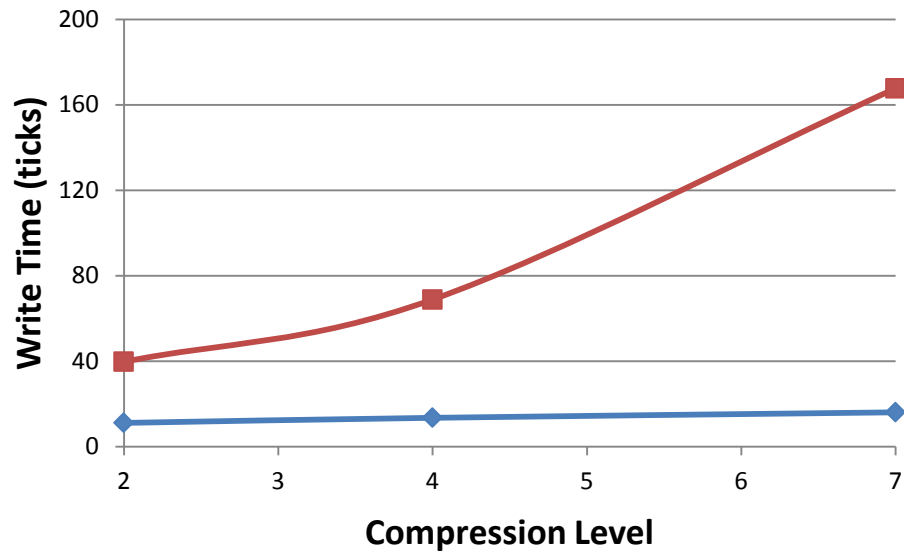


- Increasing the basket size to 262144 bytes also does not significantly increase the total RSS of the application used to create the CMS miniAOD




Example Results: Performance vs File Compression Algorithm

Write Time vs Compression Level



Output File Size vs Compression Level



- Of two algorithms, LZMA and ZLIB, the first is more complex, hence taking longer processing time, but yields smaller output size.
- Write time for LZMA rises drastically as compression level increases whereas ZLIB consumes almost the same amount of time
- Compression level has trivial effect on read time for both algorithms, however, LZMA read time could take up to 40% longer than that of ZLIB

Conclusion

- We investigated ways to improve CMS MiniAOD performance.
- We found that optimizing some parameter settings could result in significant performance gain.
- Basket size tuning could improve read back time as well as output file size of MiniAOD while having a relatively small increase in job memory
- Other parameters we investigated, including the compression algorithm used were already set close to their optimal point
- Next step: Investigate ways to improve performance by changing the miniAOD object definitions



Thank You