



# RNTuple CMS Updates

Dr Christopher D Jones

CCE SOP

23 October 2024

# RNTuple Split Leaf Fields

---

- RNTuple has a compression optimization called ‘splitting’
  - Leaf fields (e.g. UInt32 , Float64, etc) are ‘split’ by high and low bytes
  - high and low bytes are written to their own buffers and compressed separately
- ROOT team found this did not always lead to best compression
  - Based on studying files from ATLAS and CMS
- Can programmatically set a sub-field to be unsplit

```
iField.SetColumnRepresentatives ({{ROOT::Experimental::EColumnType::kUInt16}});
```

# Experimenting with Unsplit Fields

- Created MiniAOD RNTuple file with all leaf fields being unsplit
- Compared sub-field by sub-field which take less space unsplit
  - 11 selected to be unsplit
  - most important is > 10x more space saving compared to 2nd

MiniAOD Writing Summary

|                                      | File Size     | File Size Ratio | Max Allocation Difference | Max RSS Difference | Throughput 1 thread | Throughput 4 threads | Throughput 8 threads |
|--------------------------------------|---------------|-----------------|---------------------------|--------------------|---------------------|----------------------|----------------------|
| TTree                                | 4,685,680,260 | 1.000           | 0                         | 0                  | 36.98               | 71.69                | 71.83                |
| RNTuple streamer                     | 4,389,401,160 | <b>0.937</b>    | 671,073,224               | 645,521,408        | 46.42               | 160.95               | 202.51               |
| RNTuple streamer NoBufferedWrite     | 4,458,400,816 | 0.951           | -483,245,408              | -399,437,824       | 43.26               | 41.37                | 41.46                |
| RNTuple streamer zipCluster 25M      | 4,371,796,184 | <b>0.933</b>    | <b>156,402,944</b>        | <b>166,908,559</b> | 44.52               | 123.46               | 164.07               |
| RNTuple selectUnsplit                | 4,329,154,343 | <b>0.924</b>    | 706,087,944               | 668,169,830        | 46.54               | 156.03               | 183.89               |
| RNTuple selectUnsplit zipCluster 25M | 4,299,318,060 | <b>0.918</b>    | <b>173,786,808</b>        | <b>109,907,968</b> | <b>44.12</b>        | <b>135.48</b>        | <b>175.41</b>        |

# Throughput

## OutputModule Throughput vs Threads

