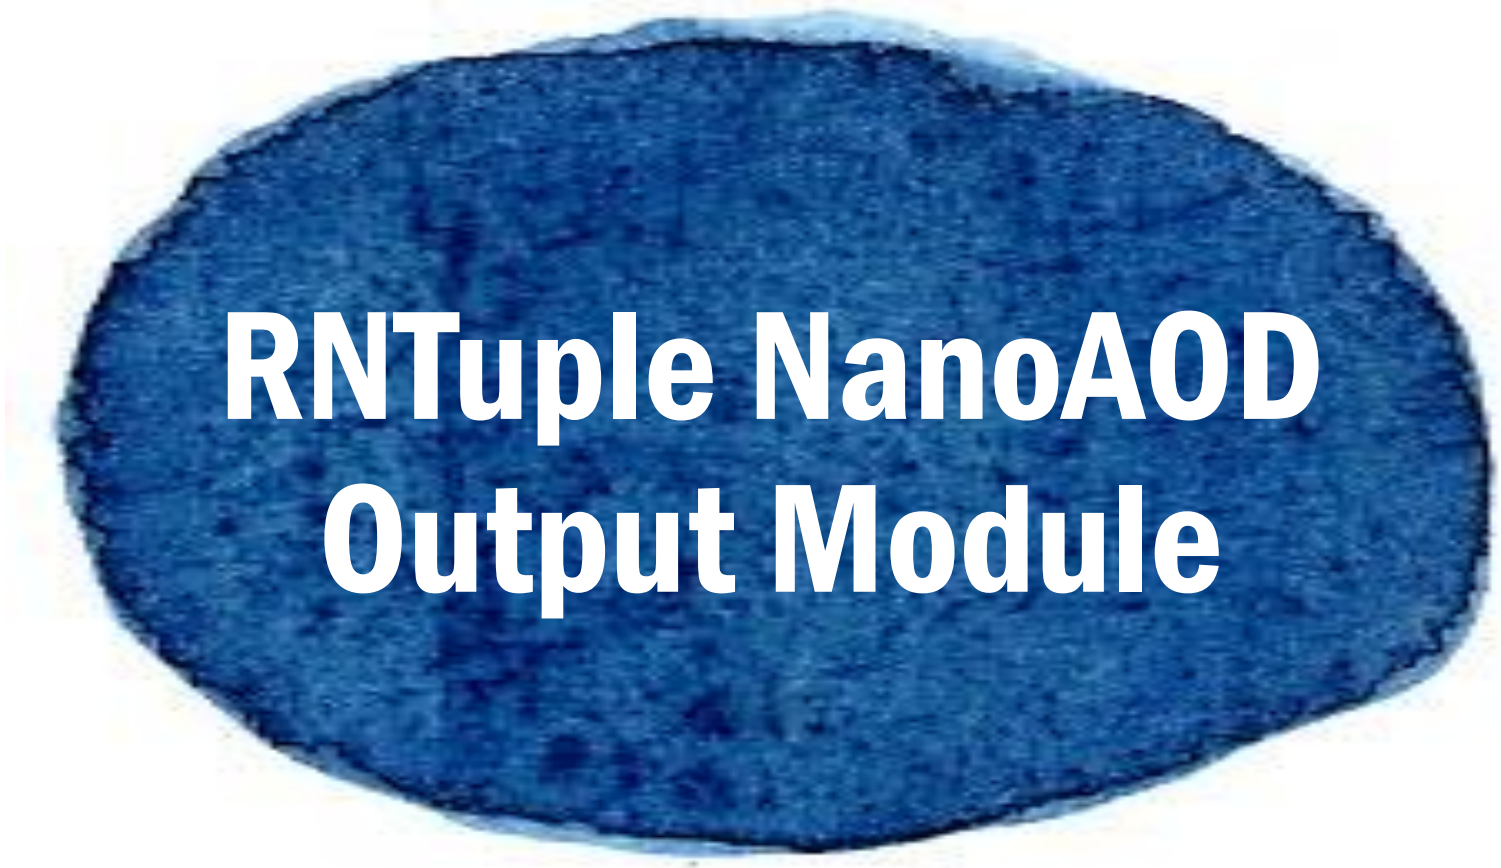

# RNTuple NanoAOD Output Module

**Max Orok, IRIS-HEP Fellow**
**Supervised by Jakob Blomer & Dan Riley**

# Outline

1. Brief introduction to RNTuple
2. Overview of prototype RNTuple NanoAOD output module
3. NanoAOD file size comparisons
4. Sample read performance benchmarks (HDD, SSD, Ceph)
5. Next steps for this project

# RNTuple

- RNTuple is the experimental evolution of ROOT's I/O subsystem
- Informed by TTree's 25+ years of production experience
- Modern implementation as a base for future improvements (e.g. first-class object store support)
- Main goals are improved performance and robustness
- There is a much more detailed overview here:

  *Evolution of the ROOT Tree I/O* (2019)

# CMS NanoAOD

- Our project goal is to develop an NanoAOD output module using RNTuple to allow for a large-scale, realistic comparison between TTree and RNTuple
- We have developed a module capable of writing the main “Events” NanoAOD tree (*backfill columns are not supported yet)
- We have verified the implementation using an automated comparator tool to ensure the branch data is equivalent
- CMSSW pull request #33825

# Storage comparison

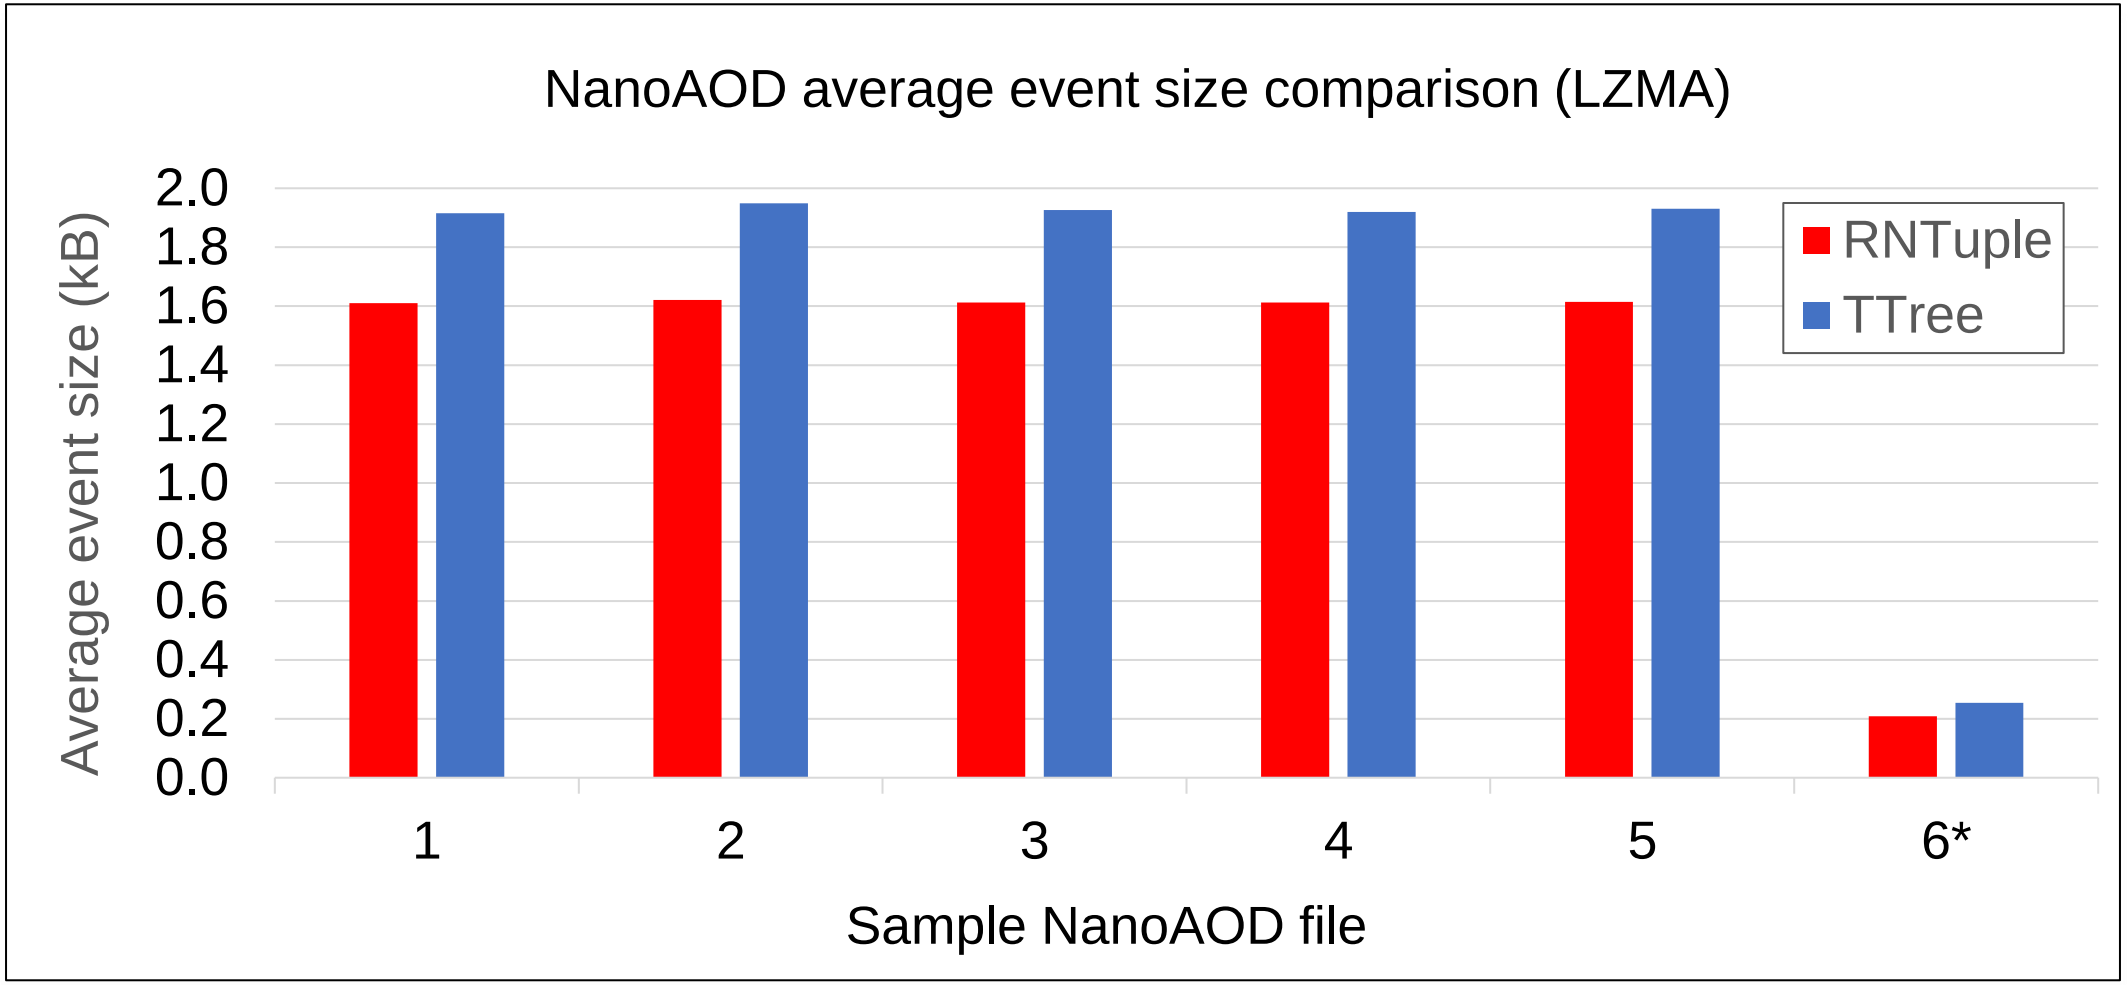


* Sample is real data, others are Monte Carlo simulations

# Storage comparison

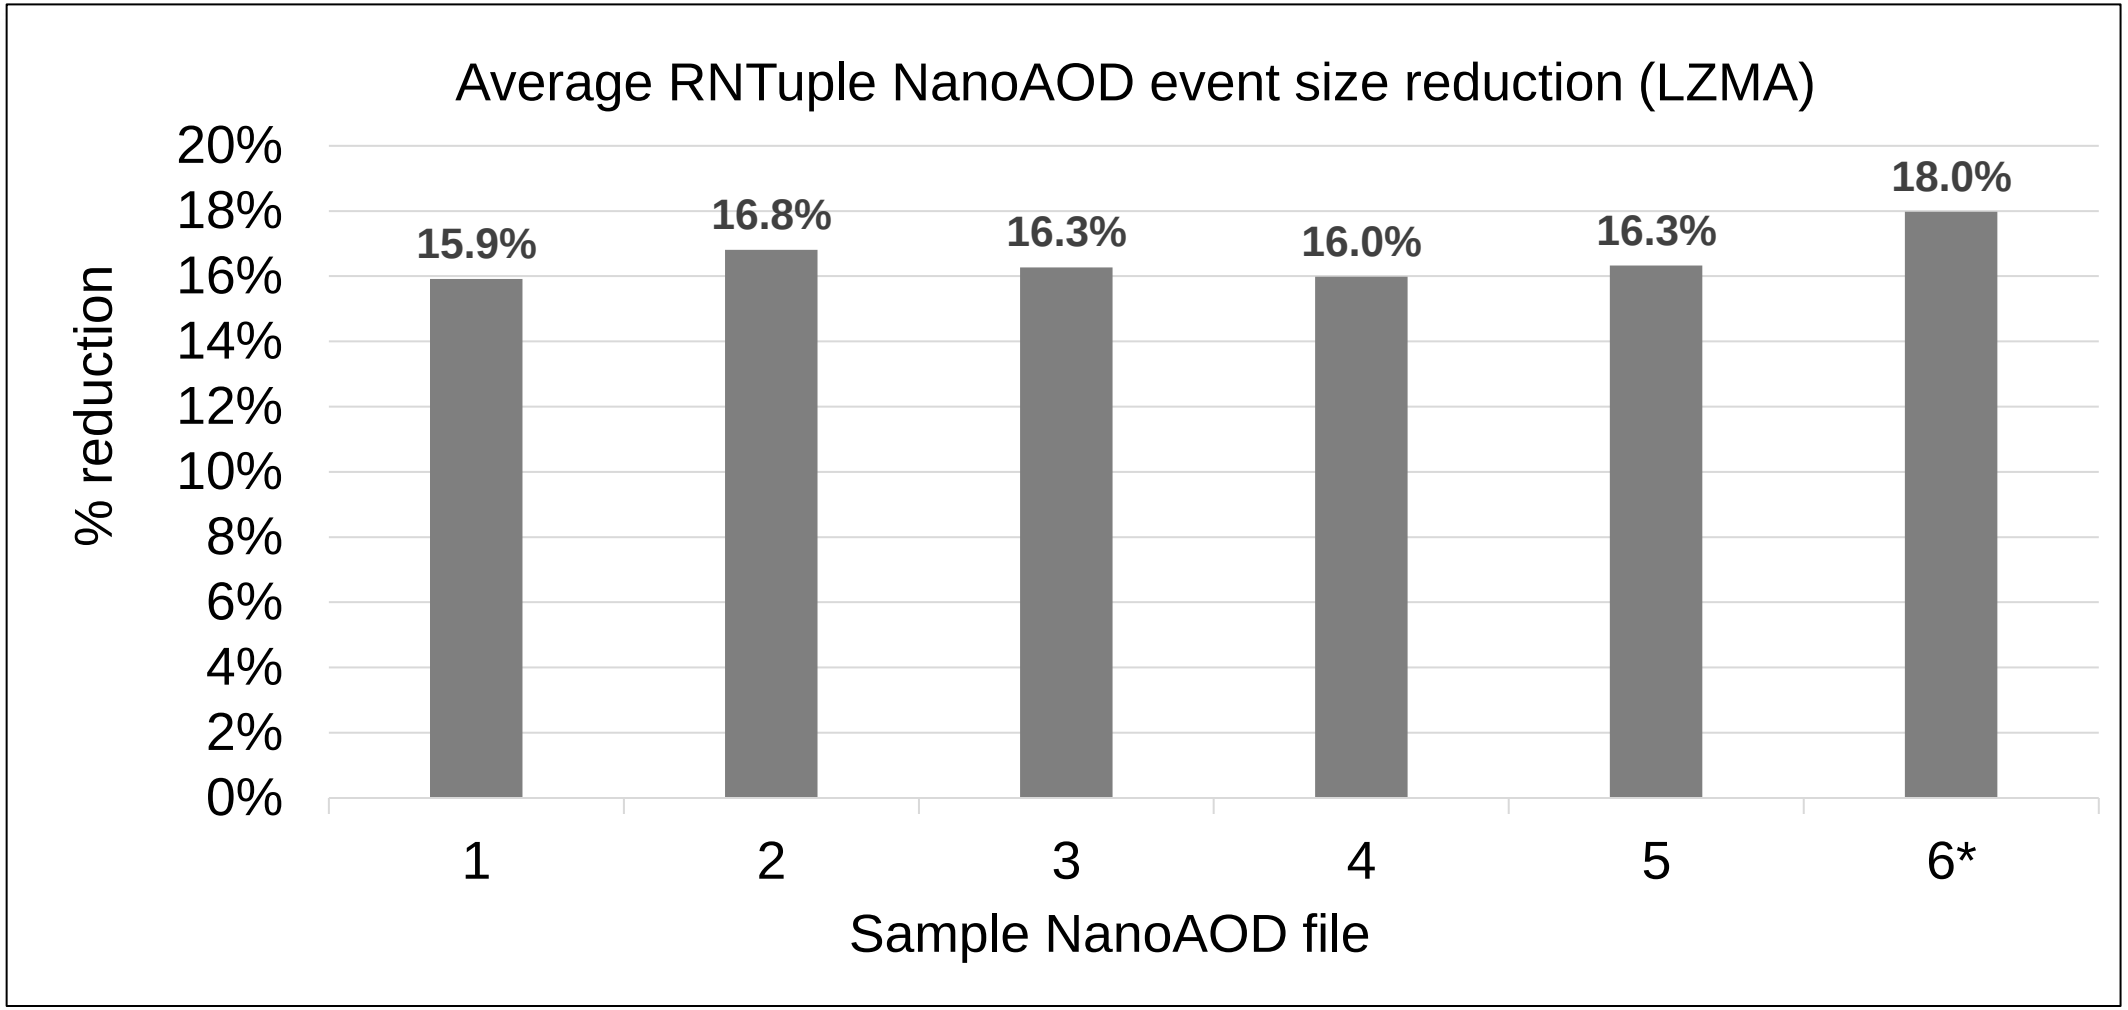


* Sample is real data, others are Monte Carlo simulations

# Read benchmark

CMS 2019 NanoAOD Dimuon spectrum

- 1.6M events, 141k selected events, sparse read (6/1479 branches)
- https://github.com/jblomer/iotools#readme

AMD EPYC 64-Core Processor

Results given for SSD, HDD, and CephFS

# SSD read results

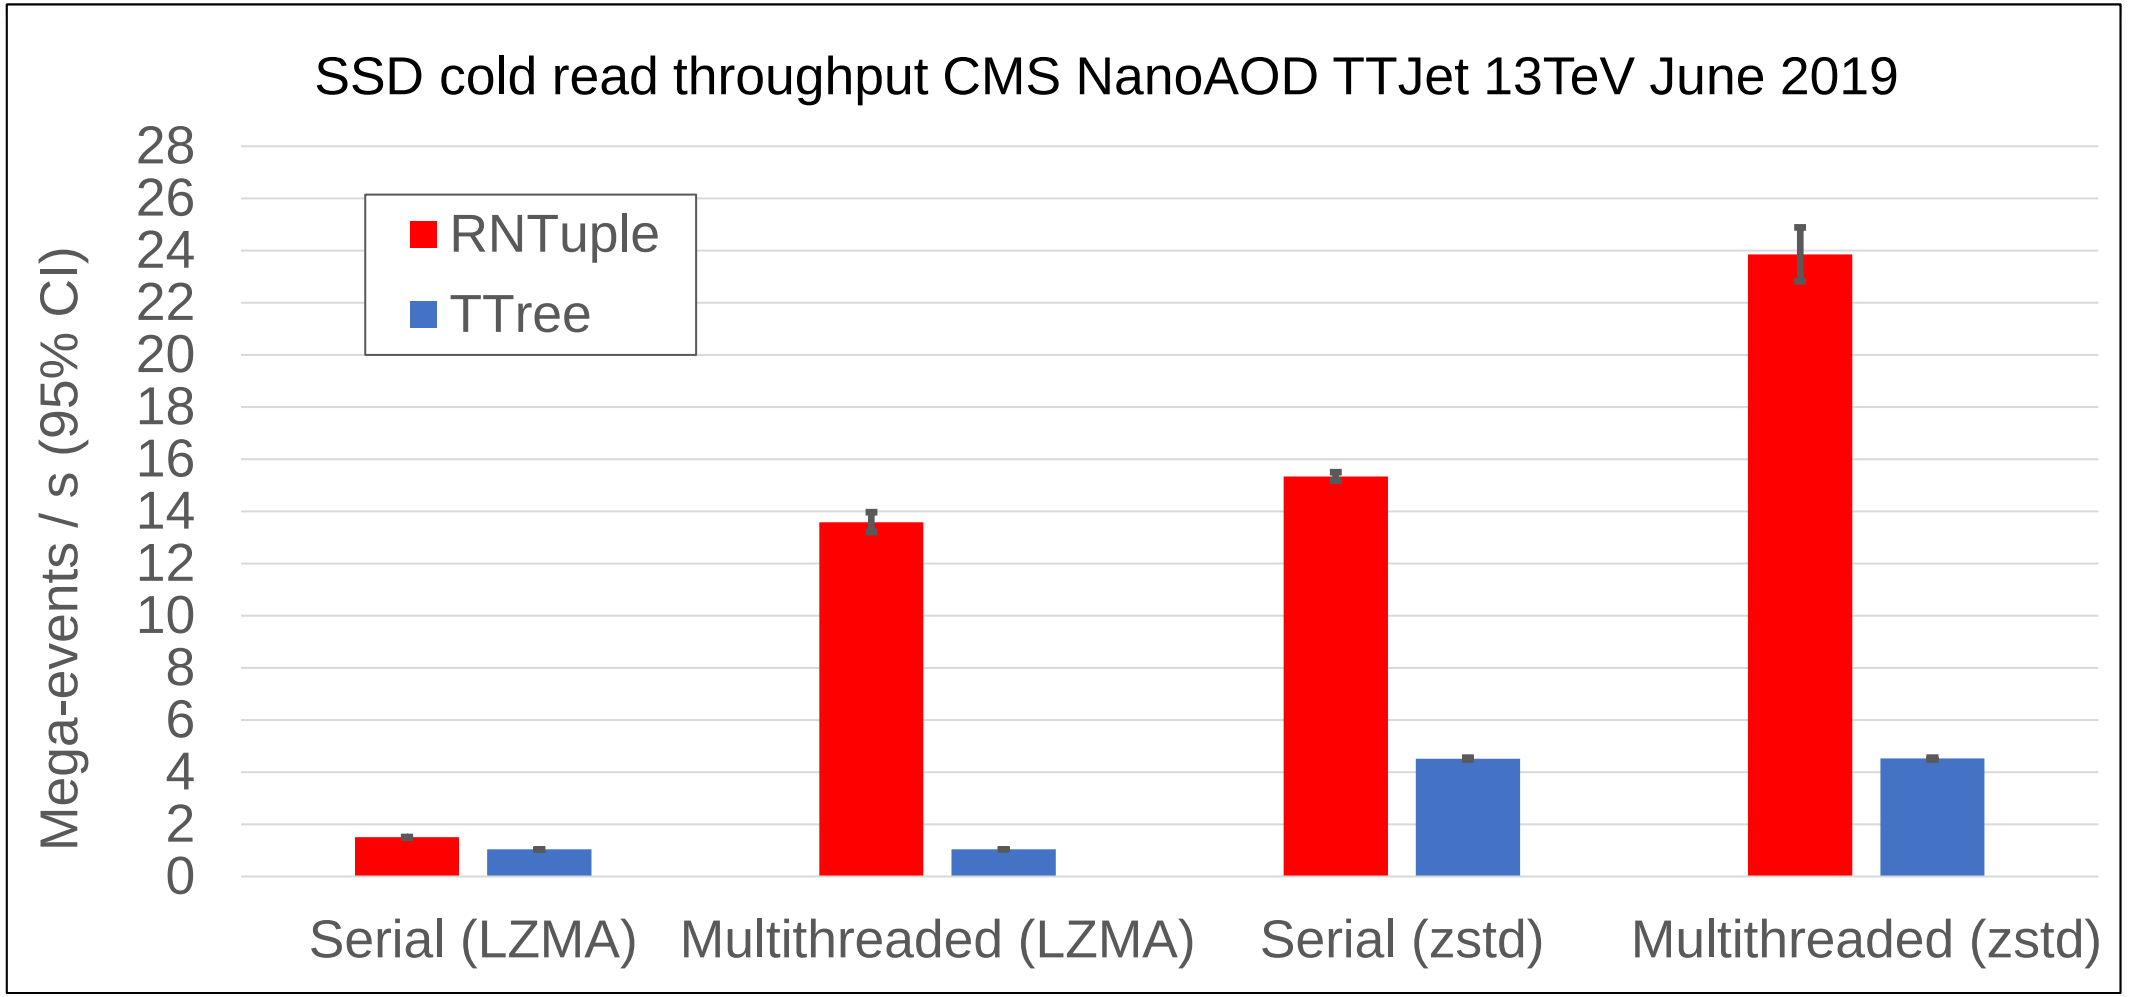

# HDD read results

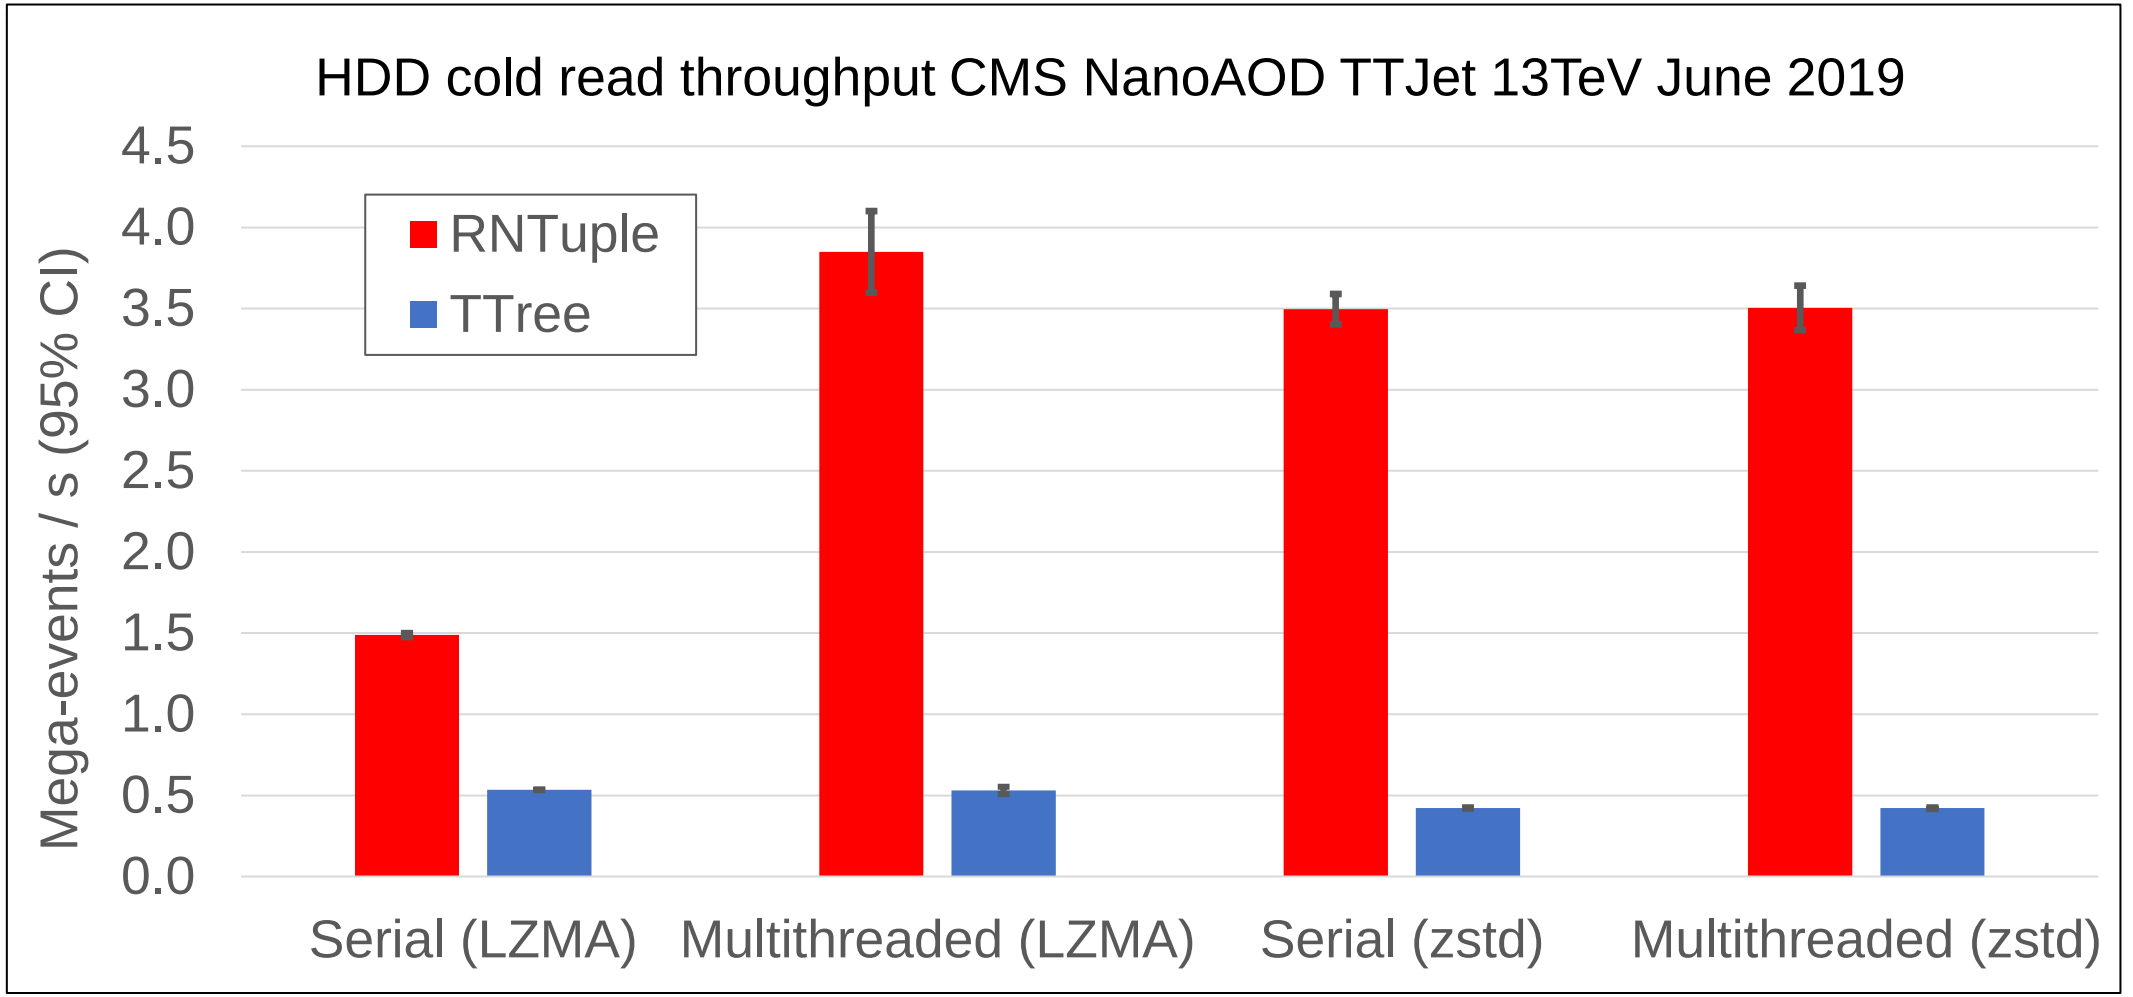

# CephFS read results

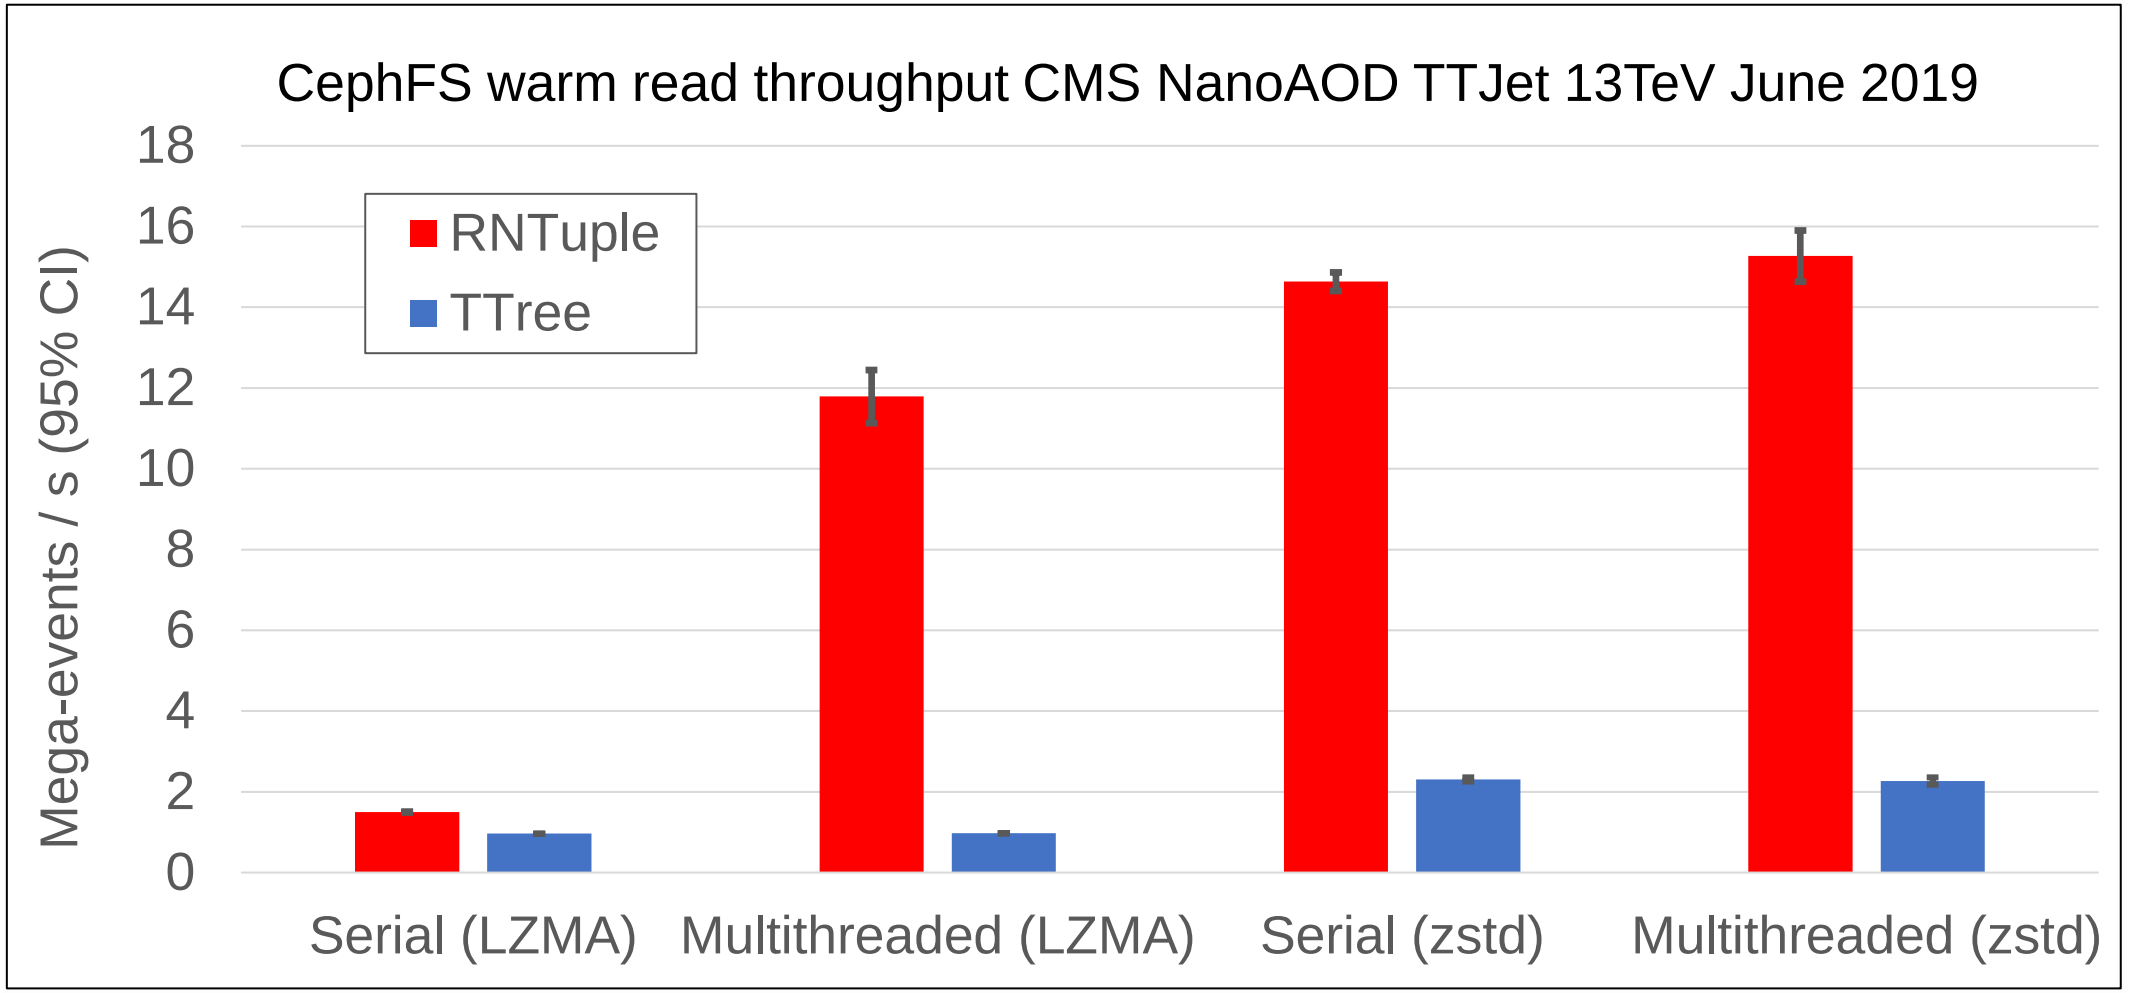

# Next steps

- Integration into CMSSW
- Backfill support through RNTuple friend mechanism
- Complete support for NanoAOD auxiliary trees
- Request for a NanoAOD open data set for performance studies and tools development

# Backup

# Page size study

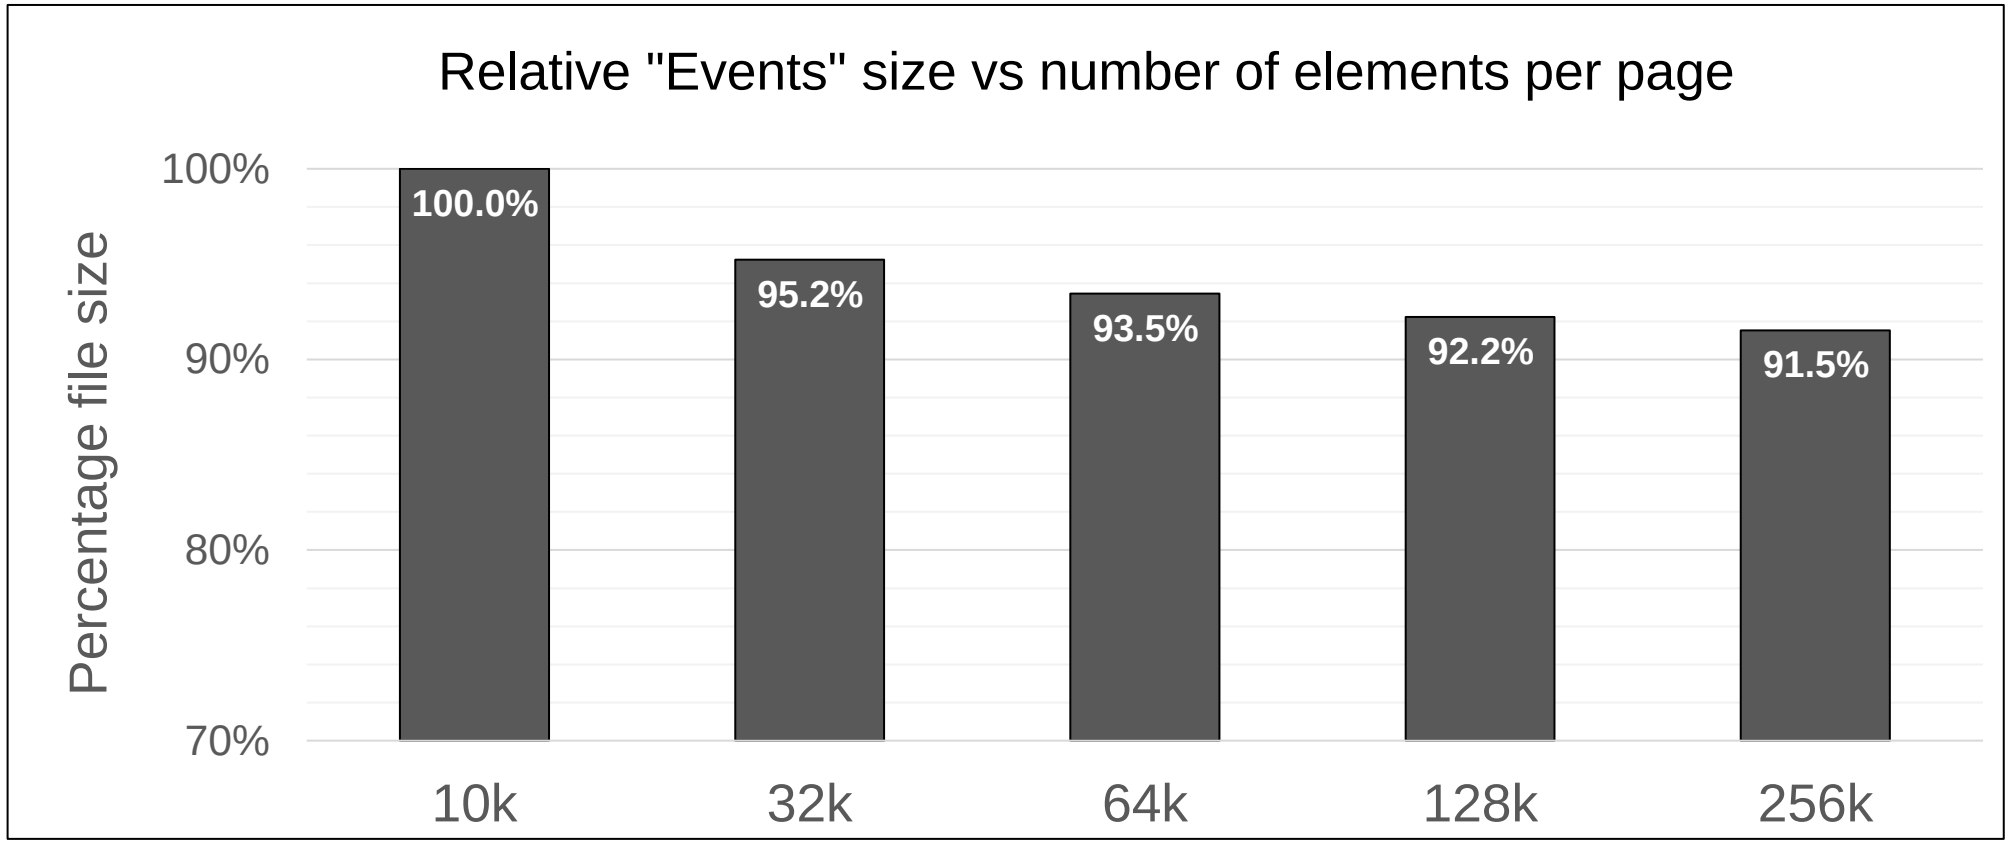


021F2782-BEB1-E611-8581-0025905A48E4.root [1]

[1] CMS Collaboration (2019). Simulated dataset QCD_Pt_1400to1800_TuneCUETP8M1_13TeV_pythia8 in MINIAODSIM format for 2016 collision data. CERN Open Data Portal.
DOI:10.7483/OPENDATA.CMS.WU0S.ZI5H

# Sample NanoAODs

Files 1–5 are from [1] (Monte Carlo)
File 6 is from [2] (Real data)

1. 021F2782-BEB1-E611-8581-0025905A48E4.root
2. 90574A80-CCB1-E611-964C-0CC47A4D764A.root
3. 289019BB-BAB1-E611-85AE-0CC47A4D76C0.root
4. 2EC9B784-BEB1-E611-B7DE-0025905A60E0.root
5. BC17960F-BFB1-E611-9723-0CC47A4C8EEA.root

6. E409E248-A7E2-BA4D-A47A-AD7F0CFD6275.root

[1] CMS Collaboration (2019). Simulated dataset QCD_Pt_1400to1800_TuneCUETP8M1_13TeV_pythia8 in MINIAODSIM format for 2016 collision data. CERN Open Data Portal.
DOI:10.7483/OPENDATA.CMS.WU0S.ZI5H
[2] /store/data/Run2017G/DoubleMuon/MINIAOD/09Aug2019_UL2017-v1/260000/E409E248-A7E2-BA4D-A47A-AD7F0CFD6275.root