

ATLAS OpenData RNTuple Conversion

CCE-SOP meeting 30 July 2025
NPPS group meeting 17 Sep 2025

Marcin Nowak, BNL NPPS

Peter Van Gemmeren, Alaettin Serhan Mete (ANL), Tatiana Ovsiannikova (University of Washington)

ATLAS I/O Group & Core Software Group, HEP-CCE



@BrookhavenLab

Introduction

- ATLAS Open Data
 - July 2024: ATLAS released 65 TB of data for unrestricted access by any researcher
 - **Open Data**: real ATLAS Athena DAOD files - same structure (all objects/branches present) but somewhat reduced content (i.e. shorter arrays)
 - Almost 500 datasets, stored in the official ATLAS data management system Rucio and EOS
 - data15, data16, mc20 - earlier data, simpler than the current Run3 data
- Conversion of the Open Data to the RNTuple format - goals:
 - Estimate filesize reduction factor on a big and varied sample
 - Exercise Athena with RNTuple in a real production environment (Panda)
 - Deliver substantial size RNTuple DAOD sample to interested users
 - Create official RNTuple datasets that are stored and registered with the production system
- Why only “conversion”
 - ATLAS DAOD production uses the one workflow (Athena MP with Parallel Shared Writer) that is not yet supported by RNTuple)
 - Therefore, for this exercise we used 1->1 Merge: Athena SP read+write job equivalent to a production job in terms of output I/O

ATLAS Datasets Used In The RNTuple Conversion

- Input

- https://rucio-ui.cern.ch/search?pattern=opendata:data15_13TeV*
- https://rucio-ui.cern.ch/search?pattern=opendata:data16_13TeV*
- https://rucio-ui.cern.ch/search?pattern=opendata:mc20_13TeV*

- Output

- https://rucio-ui.cern.ch/did?scope=data15_13TeV&name=data15_13TeV.OpenDataRNTuples.DAOD_PHYSLITE.p6026_p6898
- https://rucio-ui.cern.ch/did?scope=data16_13TeV&name=data16_13TeV.OpenDataRNTuples.DAOD_PHYSLITE.p6026_p6898
- https://rucio-ui.cern.ch/did?scope=mc20_13TeV&name=mc20_13TeV.OpenDataRNTuples.DAOD_PHYSLITE.p6026_p6898

- Created with ROOT 6.34.04 (Athena rel25 - production)

- ROOT 6.36 had troubles reading back our DataVectors written out with 6.34

Filesize reduction


Data type	Datasets	Files	Events (Mil)	TTree Size TB	RNTuple Size TB	Size Reduction %
data15	131	10049	1727	10	4.346	56.5%
data16	186	45571	5383	38	18.629	51%
mc20	79	9831	1152	12.8	5.878	54%
TOTAL	396	65451	8264	60.8	28.853	52.5%

- Extra filesize reduction caused by:
 - Simple structure and content of the data (OpenData, DAOD)
 - More robust internal Athena file format in the newer release (DataHeader smaller)
- Run 3 data shows less filesize reduction, Run 4 expected to gain even less
- E.g.: Run 3 AOD filesize gain ~**15%**



Hand-Tuning of the Compression

Non-split encoding



Size Reduction On The Field Level - Best Cases

- Per field size change in Run3 AOD, 68GB input dataset

Field	TTree size	RNTuple size	size reduction %	size reduction KB
egammaClusters_links	526600	402389	23.60%	124210
HLT_IDTrack_JetSuper_FTFAux.	717429	589037	17.90%	128392
JetETMissNeutralParticleFlowObjectsAux.	657486	522454	20.50%	135031
CaloCalTopoClustersAux.	2565370	2416516	5.80%	148855
L1_gFexDataTowersAux.	329566	173537	47.30%	156029
PrimaryVerticesAux.	799227	641286	19.80%	157941
InDetLargeD0TrackParticlesAux.	1253621	1030016	17.80%	223605
JetETMissChargedParticleFlowObjectsAux.	1559262	1219606	21.80%	339655
HLT_IDTrack_FS_FTFAux.	4173889	3439448	17.60%	734441
InDetTrackParticlesAux.	12251268	10638273	13.20%	1612995
Total size reduction:			14.70%	8054526

Size Reduction On The Field Level - Worst Cases

Field	TTree size	RNTuple size	size reduction %	size reduction KB
CaloCalFwdTopoTowersAux.	2016053	2274207	-12.80%	-258153
AllCalo	3998721	4203603	-5.10%	-204883
CaloCalTopoClusters_links	1367797	1471995	-7.60%	-104197
TauPi0Clusters_links	997103	1096859	-10.00%	-99757
JetETMissNeutralParticleFlowObjectsAuxDyn.EM_PROBABILITY	115768	167295	-44.50%	-51527
egammaClustersAux.	249636	280273	-12.30%	-30637
HLT_FastCaloEMClustersAux.	124056	149935	-20.90%	-25879
HLT_MuonL2SASInfolOmodeAux.	100372	123821	-23.40%	-23450
HLT_TrigTauRecMerged_CaloMVAOnlyAux.	98008	120818	-23.30%	-22810
JetETMissNeutralParticleFlowObjectsAuxDyn.LAYERENERGY_EME3	68044	86569	-27.20%	-18526

- Some fields take more space in RNTuple format than they did in TTree
 - In this particular case it is 1.5GB loss hidden in the overall 8GB gain
 - If the fields at least kept their original size, the overall gain would increase to 9.5GB (2.75%) more
 - This can be achieved by disabling byte-splitting feature of RNTuple field serialization (C++ API)

PHYSLITE event size: TTree vs. RNTuple (non-split encoding)

Martin Føll's Qualification Task [SPOT presentation](#)

Takeaway

Additional $\sim 2.5 - 4.5\%$ reduction on **TTree** event size when using **non-split encoding** on a handful of specific fields (compared to split encoding).

Summary of event sizes from figure:

Process	TTree [kB/event]	RNTuple [kB/event]	Size reduction [%], TTree $\rightarrow$ RNTuple
$t\bar{t} \rightarrow \text{hadrons}$	14.03	8.82	37.16 (+2.92)
$WW \rightarrow 3l, 3\nu$	11.45	6.65	41.95 (+4.62)
$Z \rightarrow \mu^+\mu^-$	11.05	5.78	47.73 (+2.53)
$e\nu\gamma\gamma$	9.11	4.61	49.37 (+2.72)
data22 (GRL)	6.71	4.06	39.53 (+3.90)
data22 (not GRL)	5.28	2.927	44.53 (+2.88)

(+X): Increase in size reduction compared to Slide 5,
e.g. $t\bar{t} \rightarrow \text{hadrons}$: $37.16 = 34.23 + 2.92$
(0.01 difference from decimal rounding)

Run 3 MC and data

