PHYSLITE - a new reduced common data format for ATLAS

James Catmore (Oslo), Jana Schaarschmidt (Washington), Johannes Elmsheuser (BNL), Lukas Heinrich (Munich), Nils Krumnack (Iowa), Nurcan Ozturk (Arlington), Serhan Mete (Argonne)
on behalf of the ATLAS Computing Activity

Introduction

ATLAS is one of the main LHC experiments with a huge and diverse and successful physics program:

- Higgs searches and measurements
- Exotics searches
- High precision SM measurements
- B-physics and light states
- Heavy-ion physics

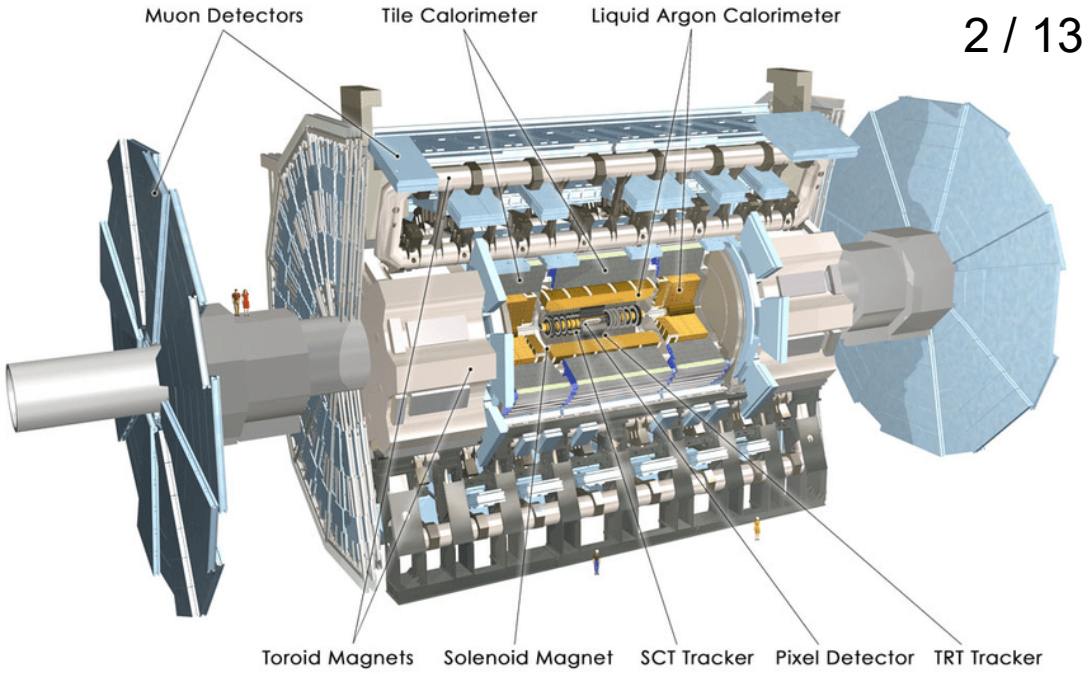

With the start of the High Lumi era in 2029, expect lots of computing challenges (increases in event sizes and rates) that demands us to improve how we use available budgets and resources. We need to reduce the needs for CPU and storage, while keep delivering high quality results in a timely manner.

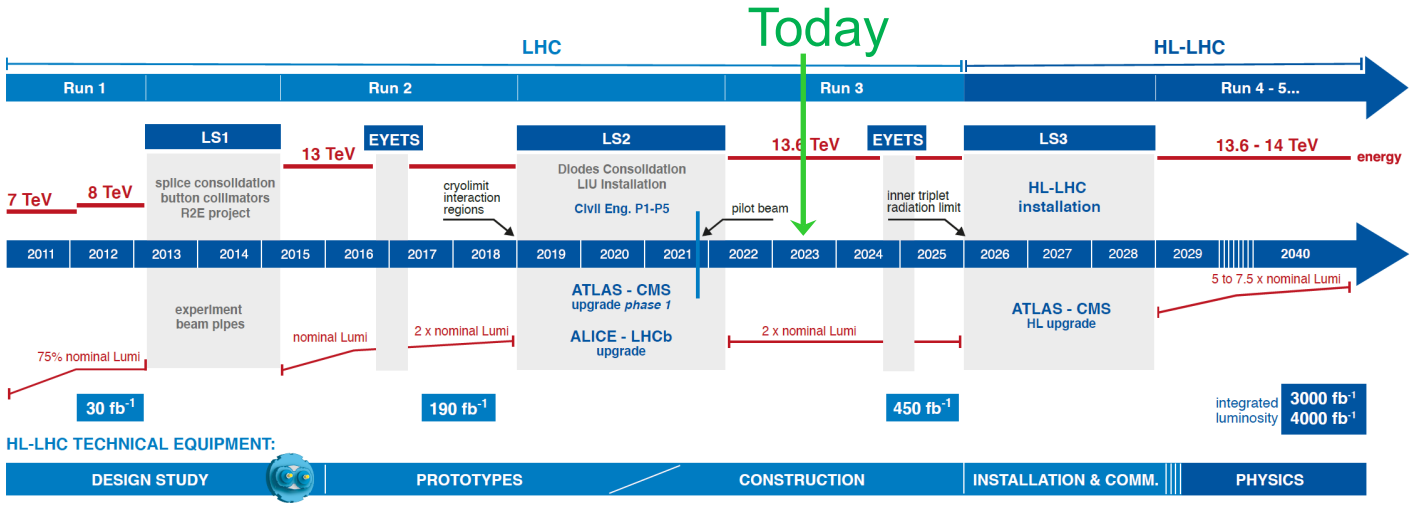

	Run 3	Run 4	Run 5
Integrated lumi / year	50-100 /fb	~270 /fb	~350 /fb
Average pile-up	60	140	200

Data Production and Reduction

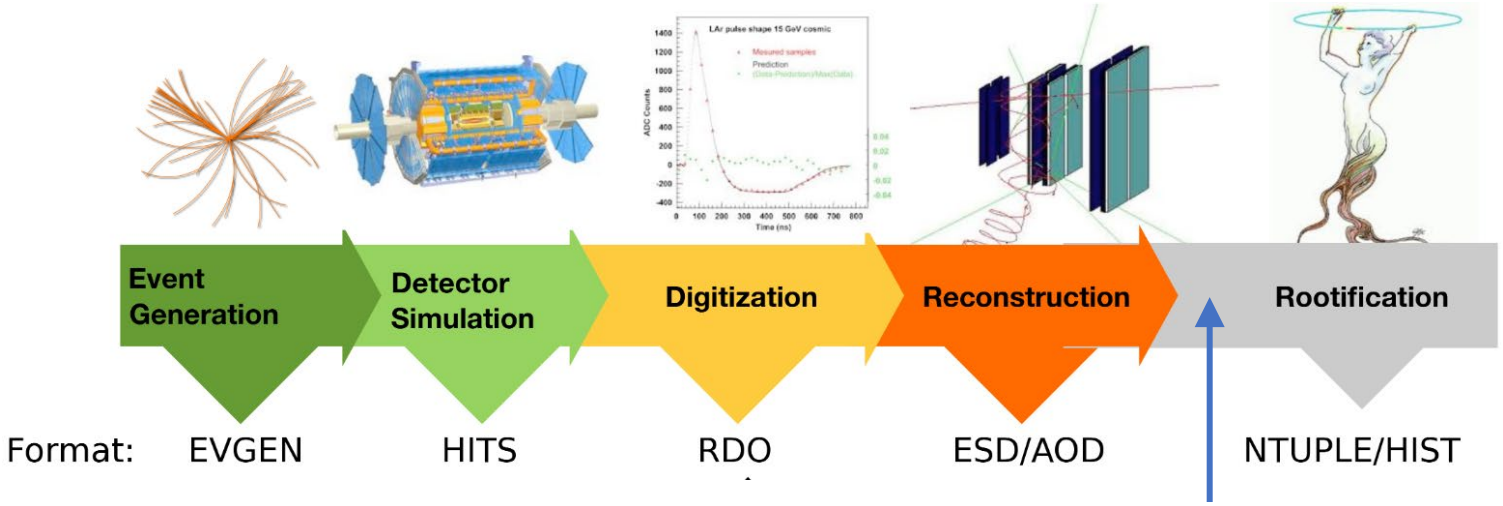

Standard AOD size: 300-500 kB/evt

Kept on tape storage but staged to disk for production of derivations („data carousel“)

Reduction of primary AOD into smaller formats called “derivations” (or short “DAOD” – derived AOD):

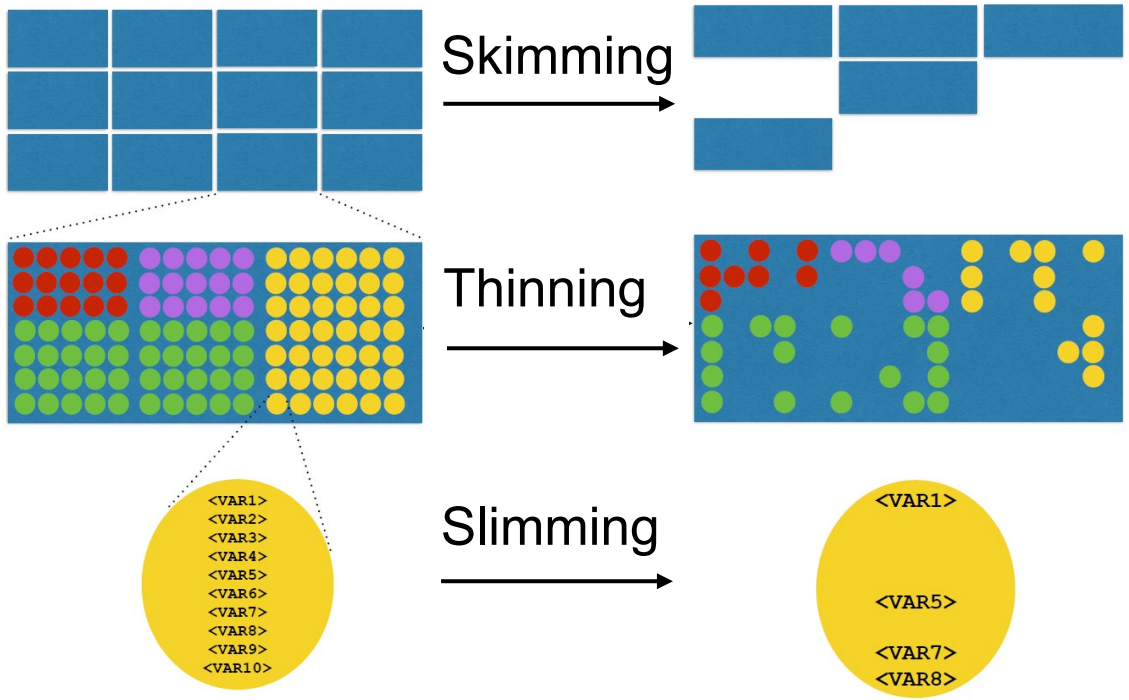

Skimming: Removal of whole events based on pre-set criteria

Thinning: Removal of whole objects based on pre-set criteria

Slimming: Removal of variables within objects uniformly across events

Derivations are reduced AODs, but often also contain additional information (eg. jet collections, high level discriminants)

Run 2 Analysis Model

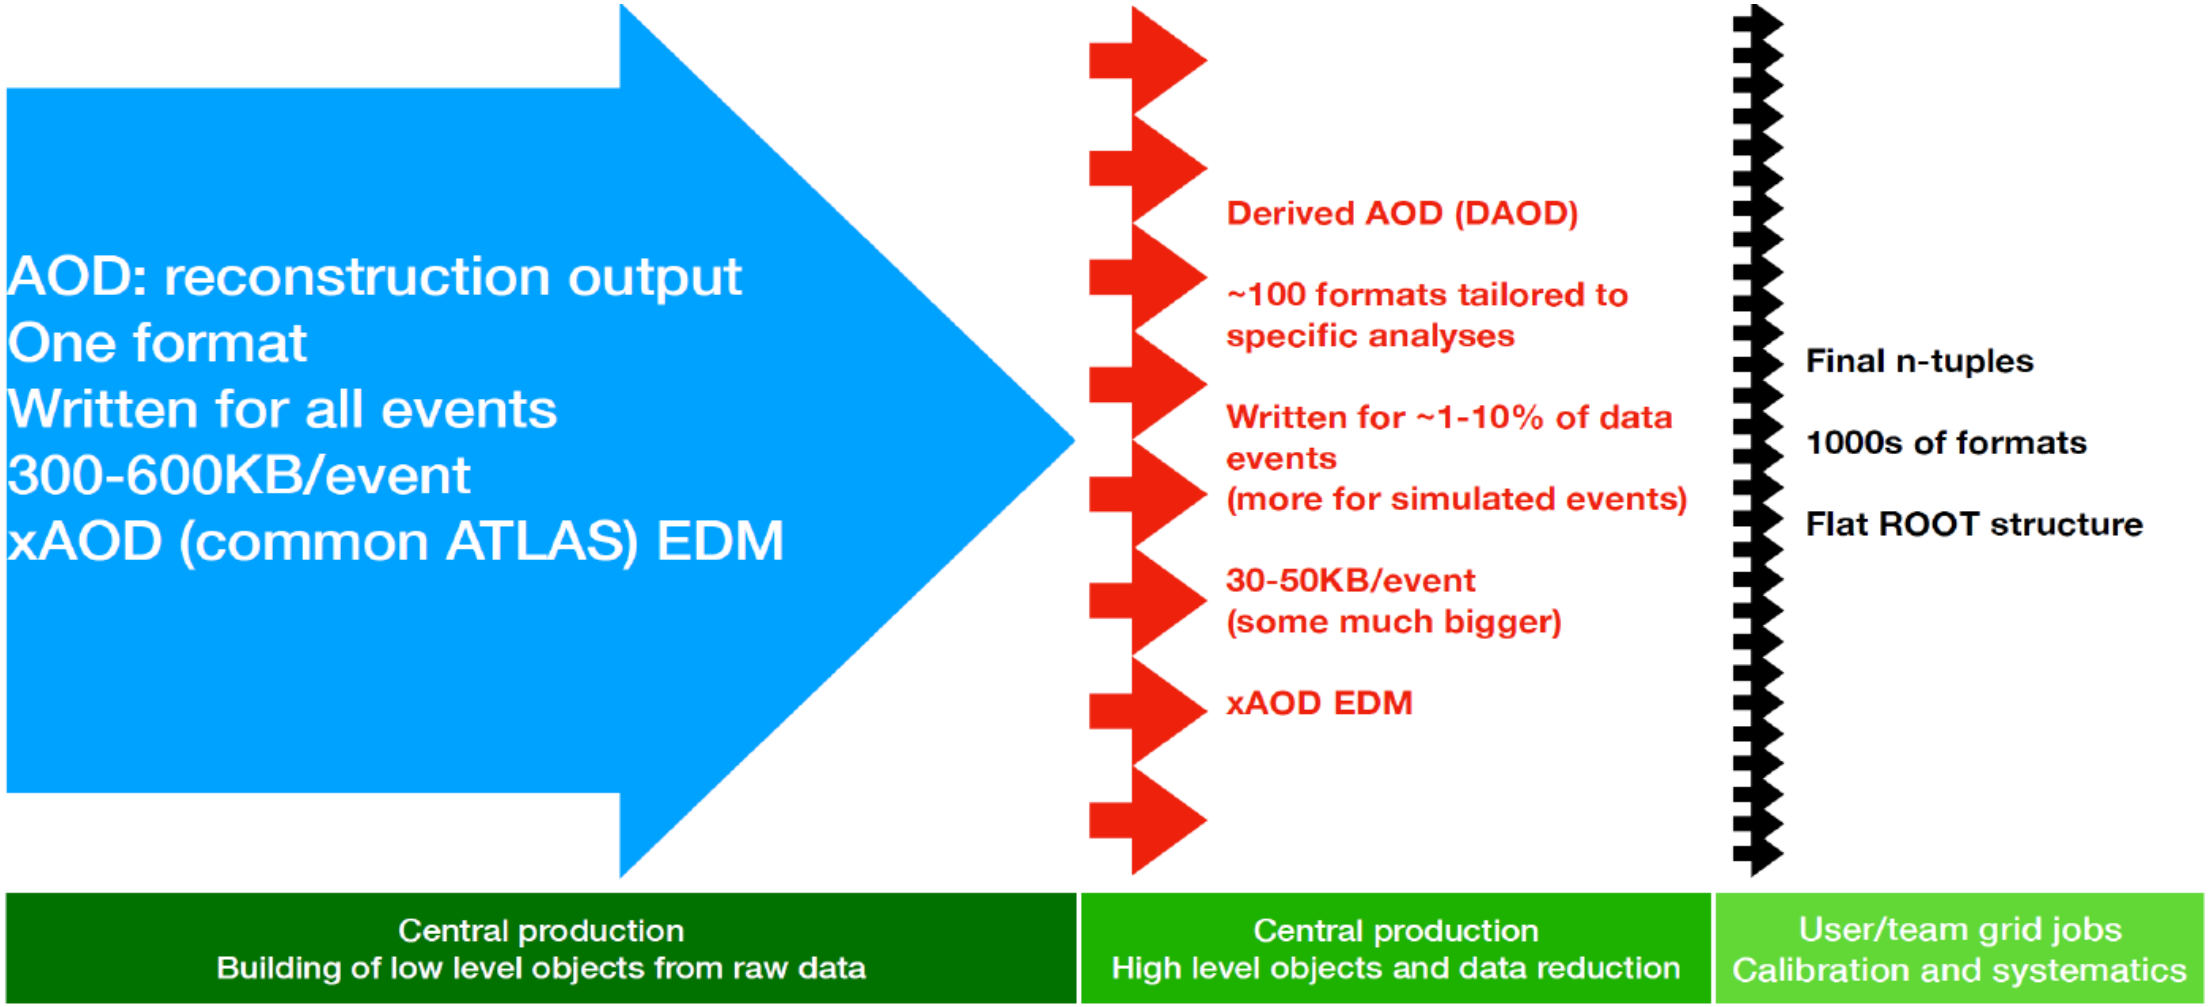

- This model worked well, and is still used to finish run-2 analyses, but it took too much disk space
- The skimmed DAOD had a lot of overlap with each other, due to loose criteria and high skim fractions
- This is inefficient especially when considering high statistics background samples (eg. V+jets, ...)
- Many formats had very similar content, meaning the same variables were stored for standard objects

Run 3 Analysis Model

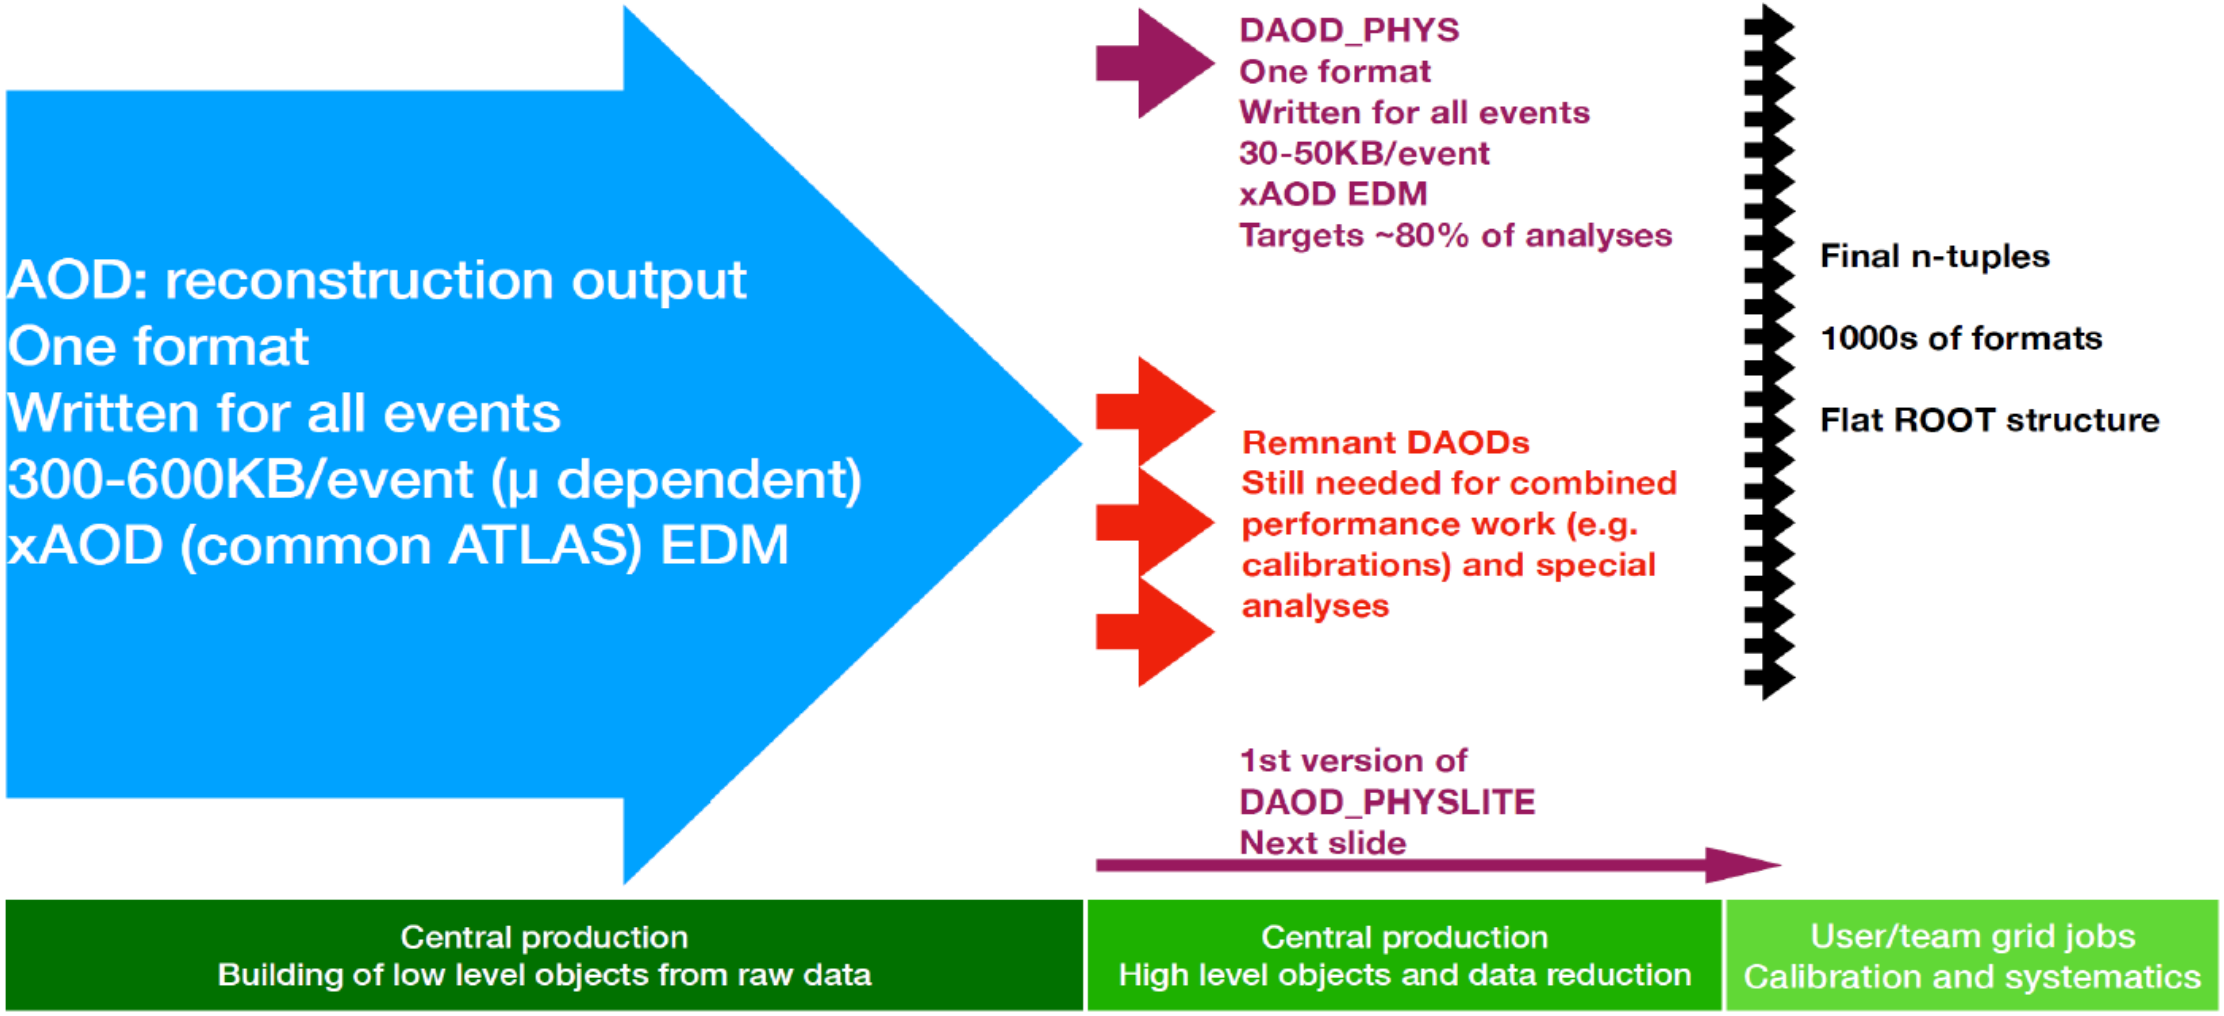

- **DAOD_PHYS** is the **common, "monolithic" derivation format** intended for ~80% of all physics analyses, not including non-standard analyses (eg. long-lived particles/signatures, custom jet collections)
- DAOD_PHYS is unskimmed, but skims can be requested (useful if efficiency is small, <1%)

Run 4 Analysis Model

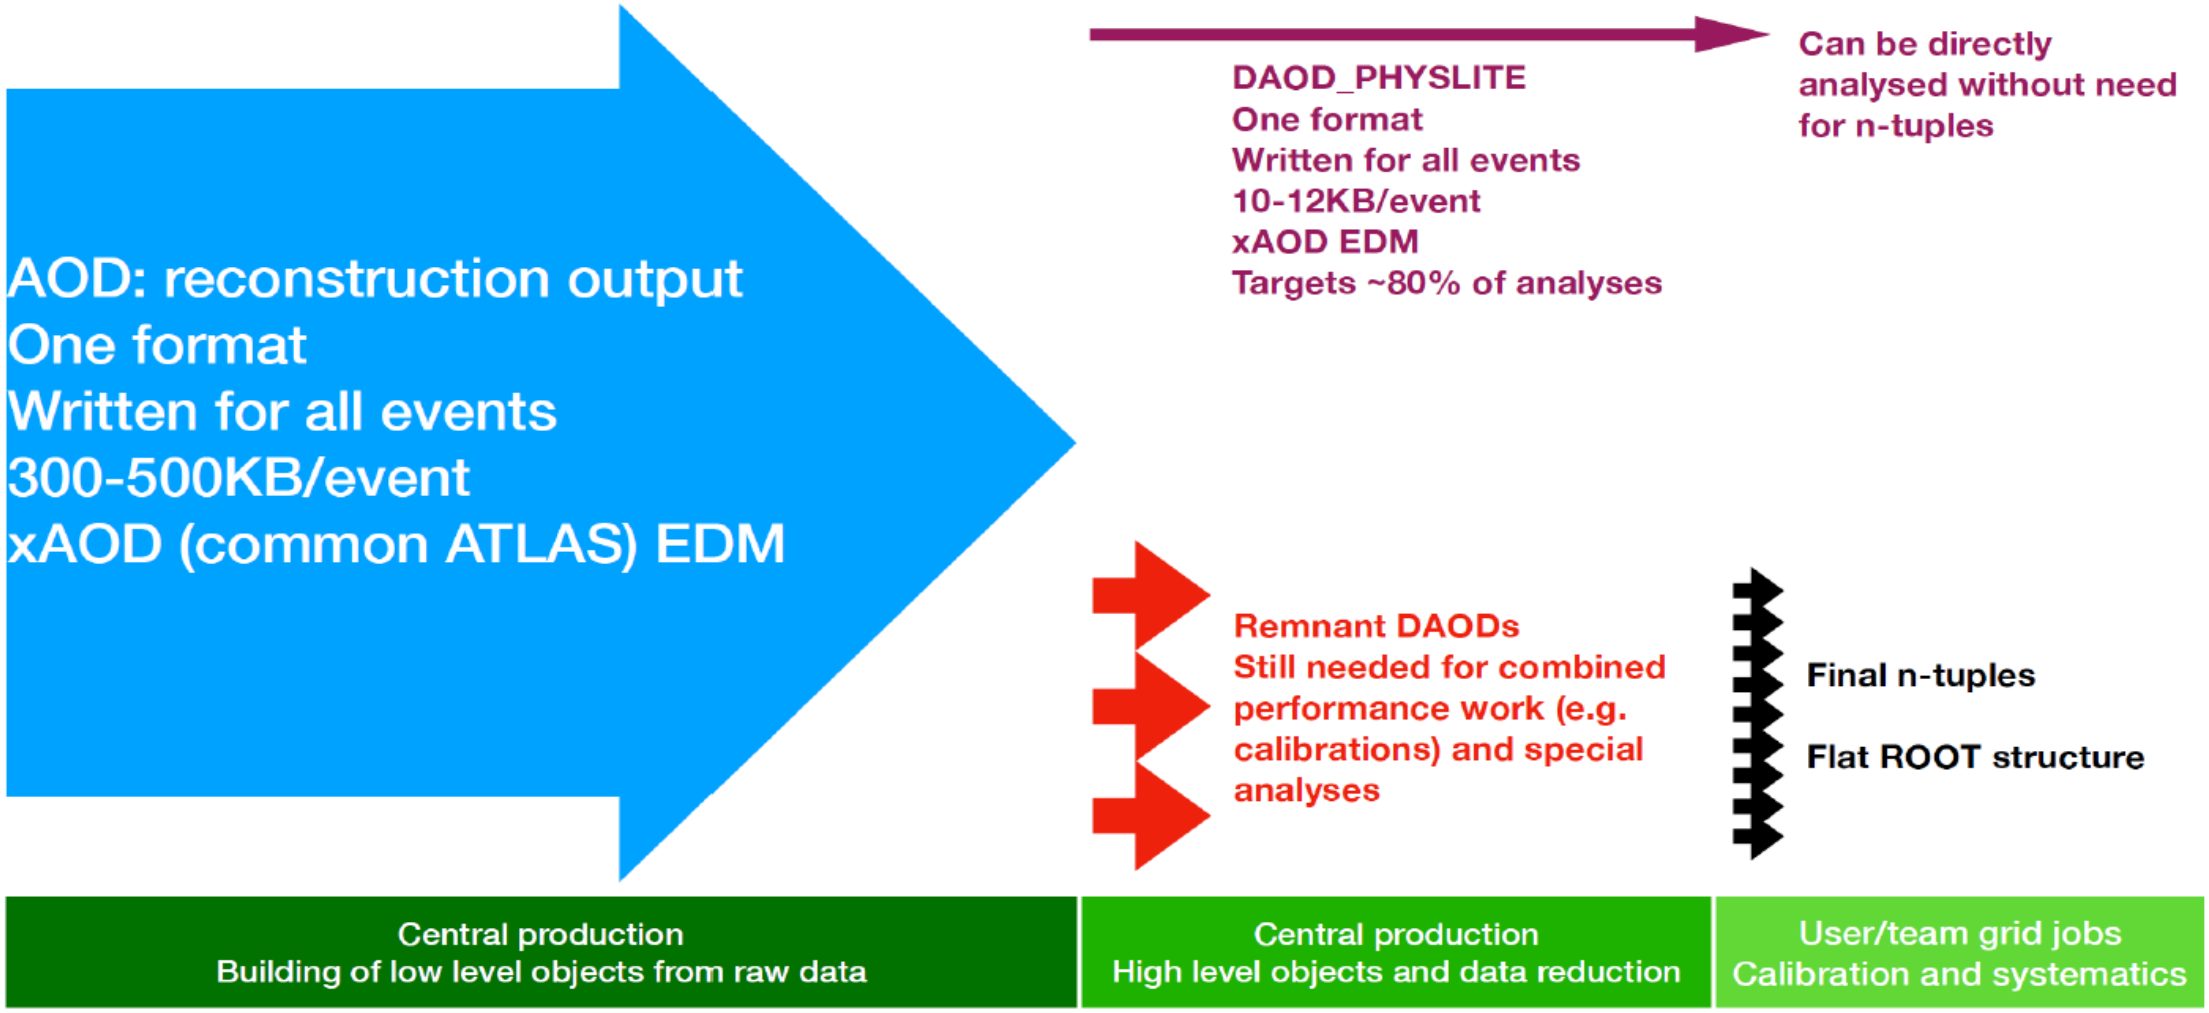

- **DAOD_PHYSLITE** is the **future format for Run 4**, user-friendly, **contains already-calibrated** objects for fast analysis. Produced by running the common CP algorithms on the raw objects from AOD or PHYS
- Like PHYS, PHYSLITE is „monolithic“, ie. one version to serve 80% of all physics analysis in Run 4.

PHYSLITE: Current Numbers

File sizes ([kB per event], using the current Run 3 prototype):

Actual size	Run 2 MC $t\bar{t}$	Run 3 MC $t\bar{t}$	data17
PHYS	34.2	40.7	21.7
PHYSLITE	13.6	16.3	6.2

Target size	MC	Data
PHYS	50	30
PHYSLITE	12	10

Work is ongoing to further reduce PHYSLITE size

Estimate for the total Run 3 dataset, assuming targeted PHYS/LITE sizes are met:

	#Events [10^9]	Size [kB/event]	#Replicas	#Versions	Total [PB]
PHYS Data	19	30	4	2	4.6
PHYS MC	24	50	4	2	9.6
PHYSLITE Data	19	10	4	2	1.5
PHYSLITE MC	24	12	4	2	2.3

Impact on CPU, storage and network:

- As PHYSLITE objects are preselected and don't need to be calibrated, can save **~25% of CPU time**
- Currently, PHYS has to be stored in addition to PHYSLITE, so no direct reduction of the total storage
- The entire PHYSLITE data production can be stored on a single grid site, which is useful to optimize network traffic (less down time, less transfers) and allows easier data access at local sites

PHYSLITE Composition

Run 3 MC (semi-leptonic tt sample)
with average pile-up of <μ> = 45

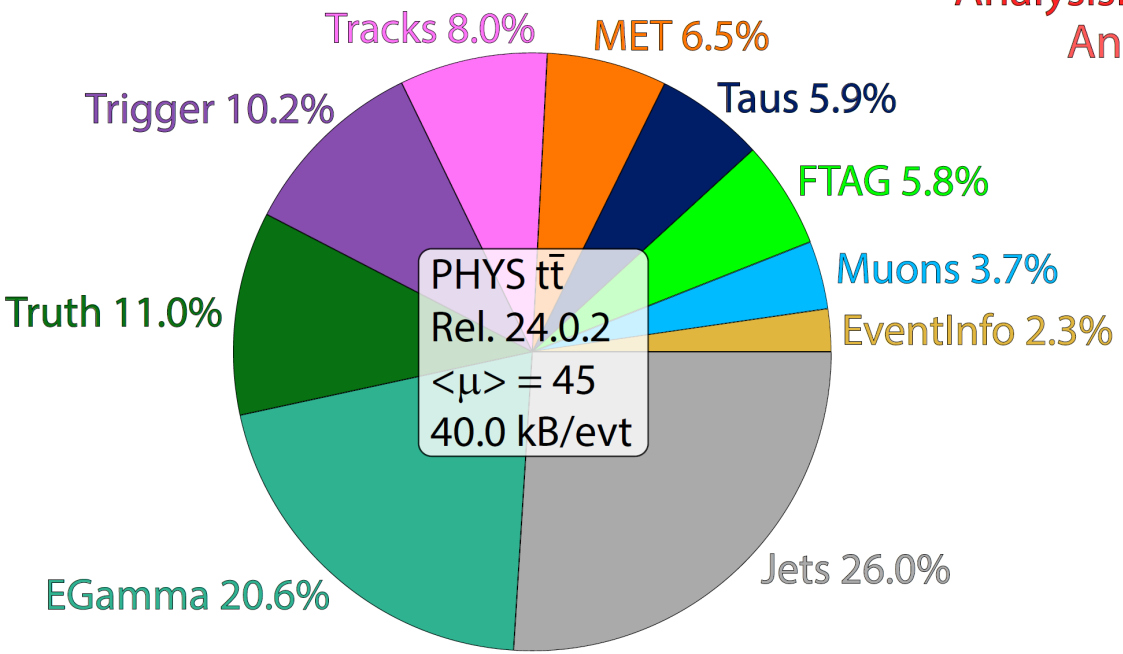

ATLAS Simulation Preliminary

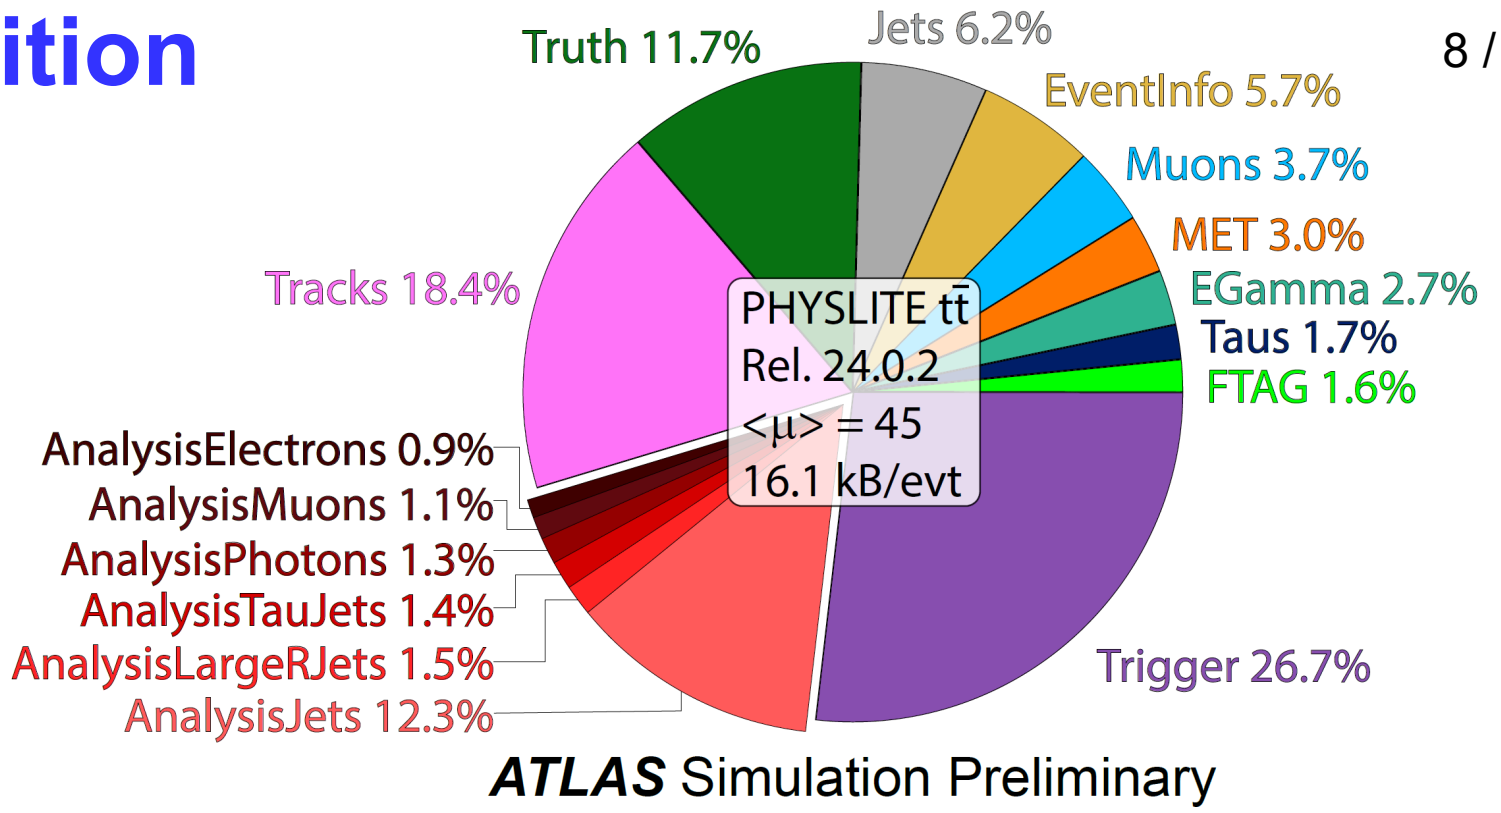

ATLAS Simulation Preliminary

In PHYSLITE, "AnalysisXYZ" are the containers holding the calibrated and preselected objects.

Trigger, tracks and truth information are the largest containers. We plan to reduce at least the tracks.

All other information is needed to run systematic variations using performance tools (C++ based).

PHYSLITE Validation

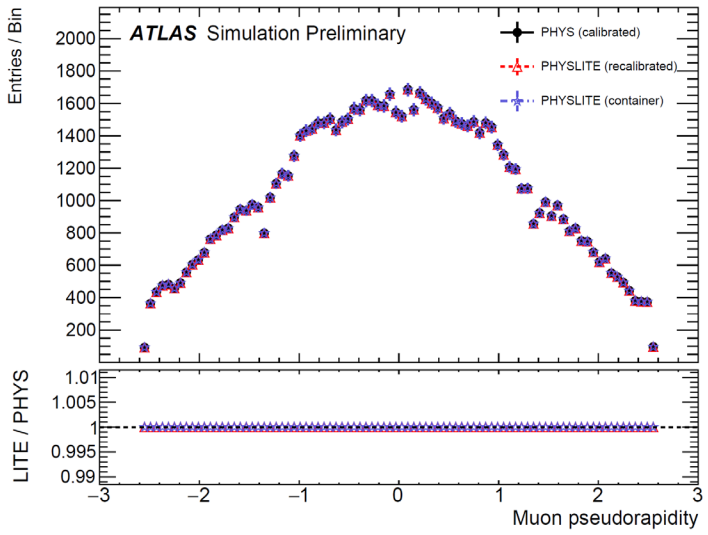

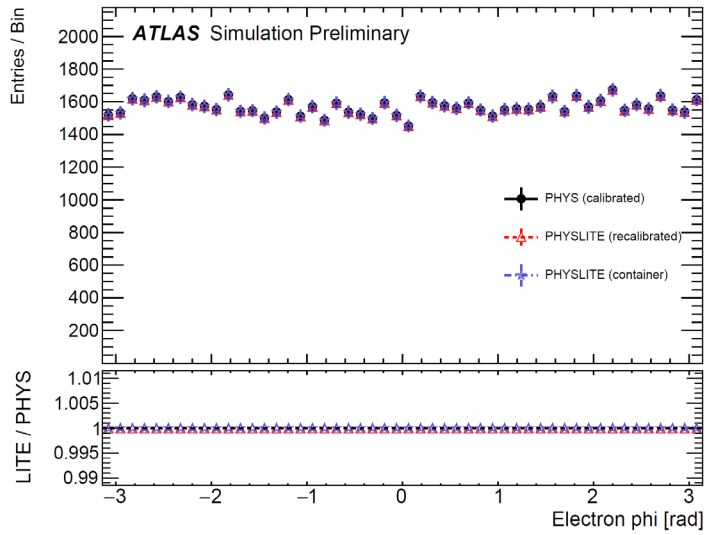

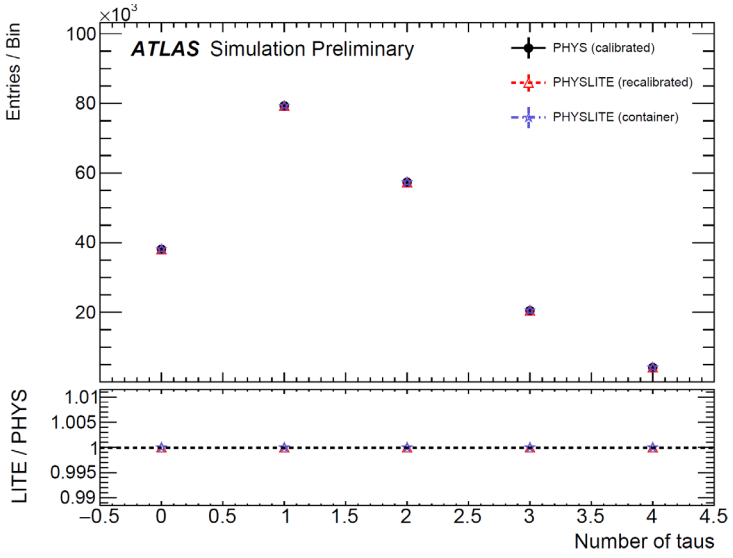

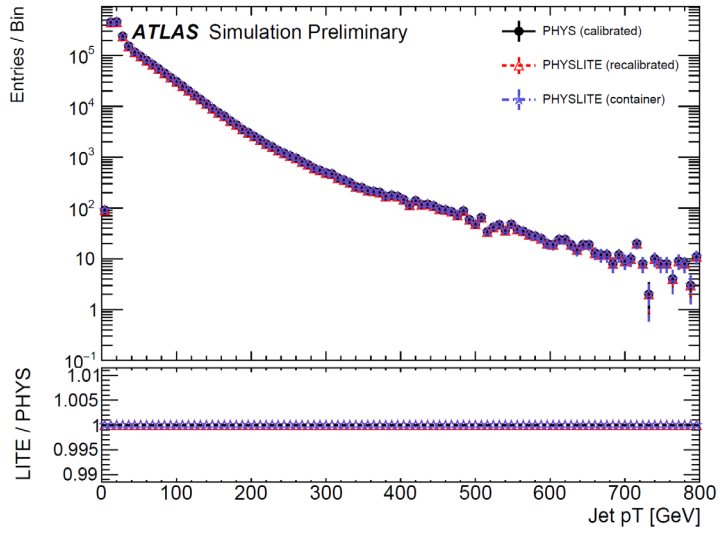

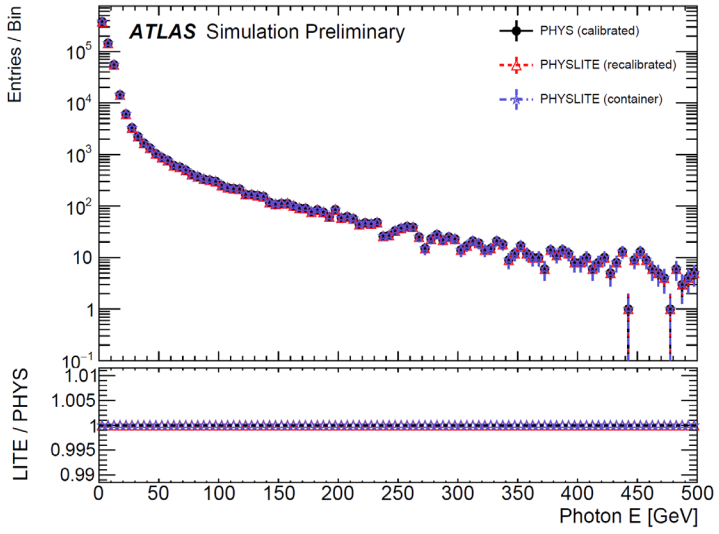

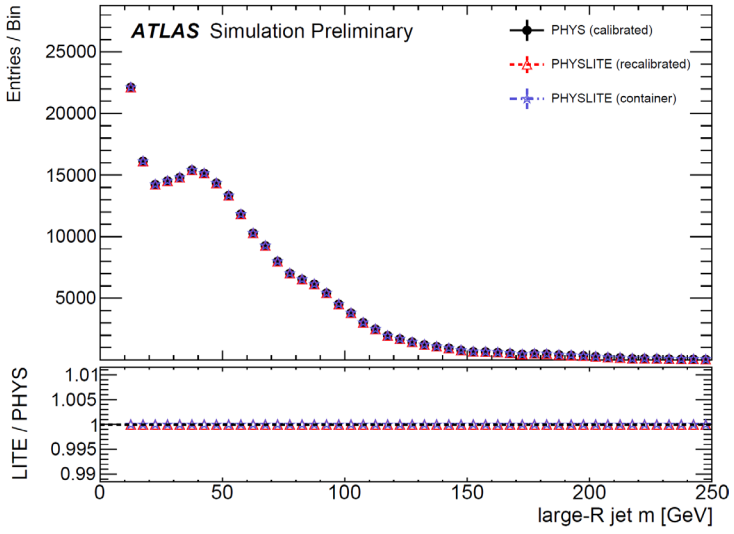

Perfect agreement in nominal object multiplicities and kinematics. Applying calibration again does not change the objects. Systematic variations have also been checked -> also perfect agreement!

PHYSLITE Production

Production schedule:

- AODs are stored on tape and will be staged at most 4 times per year to produce PHYS/LITE and residual (special) DAODs
- In the (near) future PHYSLITE will be produced from PHYS, this allows for more frequent updates (up to 6-8 times per year)
- In general, we will reprocess PHYSLITE when performance recommendations (ie. calibrations) change

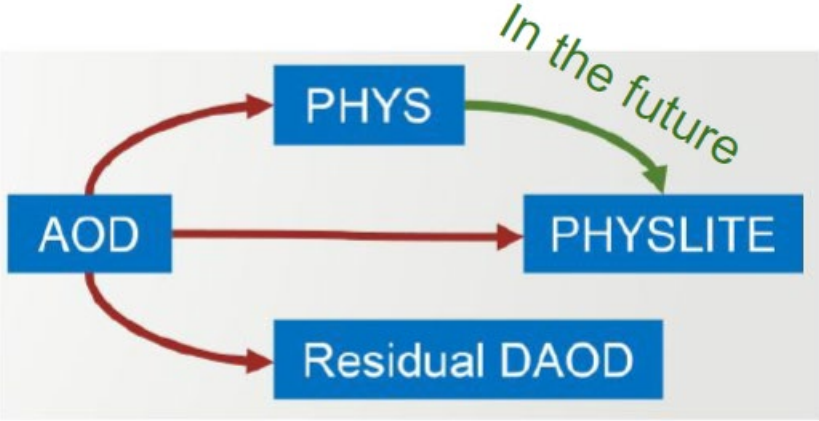

Total derivations size:

- In Run 2, all DAODs together made up 100-120% of the size of all AODs
- In Run 3, all DAODs together will amount to 50-70% of the size of all AODs

PHYSLITE is not a fixed format!

- It is meant to change and evolve with analysers' needs!
- Requests to add new features are encouraged, but impact on CPU and storage must be understood, and a set of automated validation checks must be passed

Automated Validation

- **Scripts** to check size and content for AOD and DAOD
  - Lists event sizes, container sizes, object counts, compression rates
  - Also performs integrity checks, for example checking decorations

Object counts per event follow...

NAME	Size per event (KB)	Objects per event
Kt4EMPFlowEventShape	0.009206983154557822	1.0
MET_Core_AnalysisMET	0.044394966928554574	2.0
EventInfo	0.05448895041105065	1.0
ExtrapolatedMuonTrackParticles	0.08302295415978402	0.5732985383323386
GSFConversionVertices	0.09685924562488722	1.2307852472897378
MuonSpectrometerTrackParticles	0.10667766987722269	0.9671159478630225
AnalysisElectrons	0.11616919965713969	0.18427006843325328
CombinedMuonTrackParticles	0.11859651481014284	0.5168445746172319
AnalysisTauJets	0.12071173825928985	0.729043183742591
AnalysisPhotons	0.13123319912879472	1.401785094598909
egammaClusters	0.14964488808890075	1.5860551630321622
BTagging_AntiKt4EMPFlow	0.18979453757235462	6.846601240960009
AnalysisMuons	0.19307965418563736	0.2430421565497425
TauTracks	0.22240376248421037	13.572257464499382
PrimaryVertices	0.36633920104957596	19.581335628322762
GSFTrackParticles	1.2217328527548548	14.06064602101581
AnalysisJets	1.3675272316718952	6.846601240960009
InDetTrackParticles	1.7524529895519565	17.935384017434515

- **CI pipeline tests:** Run performance tools on PHYS and PHYSLITE
  - Pipeline will fail if variables needed for these tools are not present
  - Produces output ntuples and checks for binary compatibility

- **ART tests and validation plots** for AOD and DAOD
  - Runs daily, as part of the Athena release testing
  - Outputs comparison plots of basic kinematics (#, pt, eta, phi) overlayed to a pre-defined reference

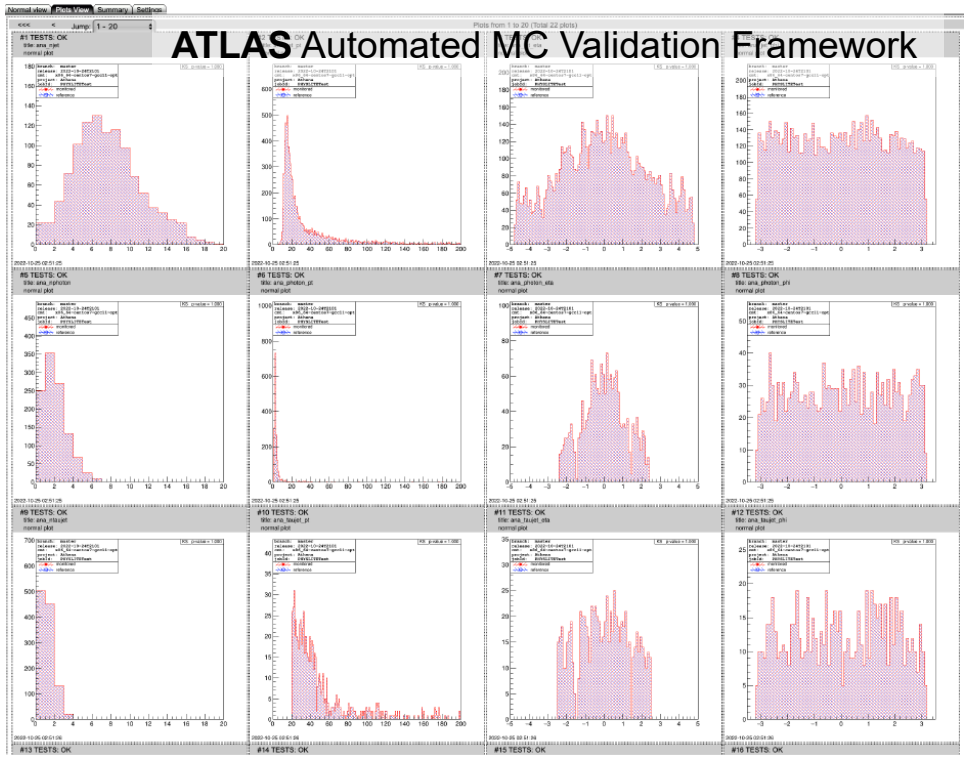

- **SPOT** (Software Performance Optimization Team) monitoring:
  - Makes plots of CPU utilization, memory allocations
  - Also check container sizes

- These automatized checks are complemented by manual validations using physics analysis frameworks

Ongoing Developments for PHYSLITE

2022/3: Run 3 analysers use **PHYSLITE prototype** and give feedback
(usability for physics analyses, bug fixes, optimizations, adding/removing variables)

-> We are continuing to reduce the size of PHYSLITE based on this feedback.

2023/4: Work towards **PHYSLITE Run 4 version**:

- RNTuple replaces TTree for storing data (RNTuple is faster, columnar, type-safe)
  -> was discussed eg. in this CHEP talk on Monday
- Event augmentation (adding additional (friend) trees for subset of events)
  -> see this CHEP talk earlier today
- Lossy compression (Storing floats with reduced precision. Infrastructure is in place, but decision which variables to compress and extensive validations are still needed)
  -> see this CHEP talk earlier today
- Columnar analysis (Main task: C++-based detector performance tools need modification)
  -> will be discussed in detail in the following CHEP talk by Nils Krumnack

2026: **Run 4 version:** PHYSLITE replaces PHYS, super-fast lightweight analysis now possible

Conclusions

- PHYSLITE is a fundamental element of the ATLAS S&C preparations for HL-LHC to reduce the storage resource needs
- A Run 3 version is already available and physics analysers started using it
  -> Profit from reduced CPU times, faster grid processing and reduced storage
- PHYSLITE is intended to replace PHYS in Run 4, serving at least ~80% of all physics analyses
- Several developments ongoing to elevate the current prototype to a faster and more optimized version, that can be augmented to serve also non-standard workflows

Backup

Storage Projections

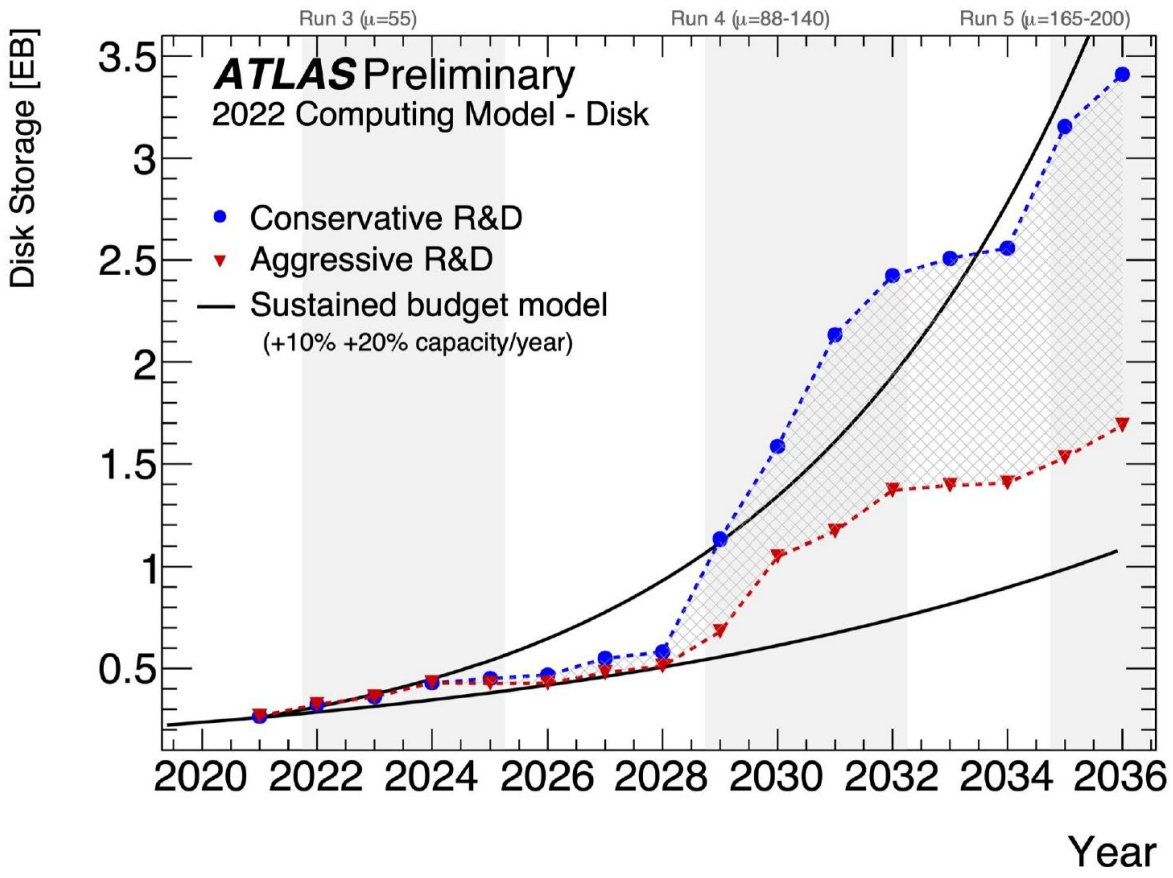

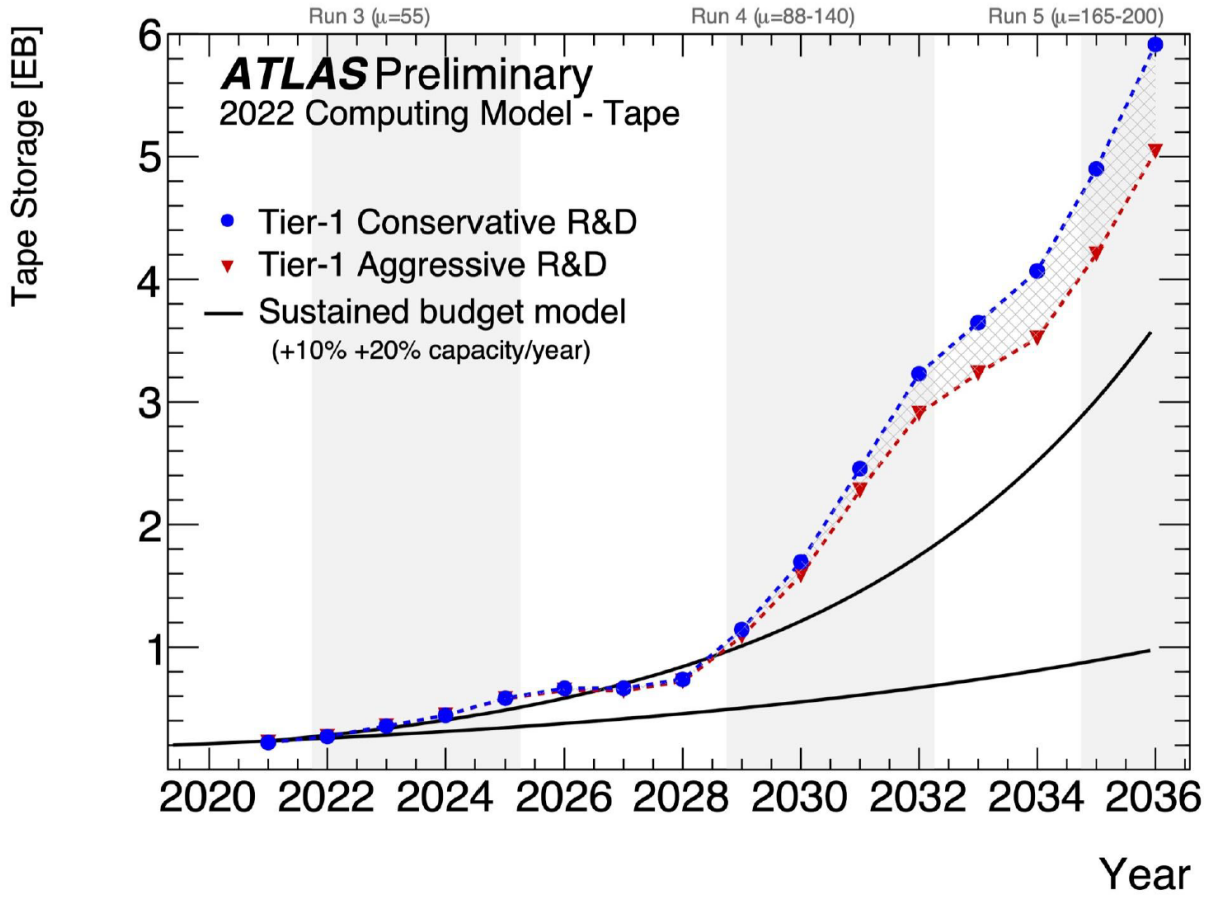

- Goal is to store as much as possible on tape.
  Tape is cheaper than disk, but data access is slower (-> data carousel to rotate information on disk)
- **Key ideas:** Multi-level data reduction resulting in small, commonly used analysis formats
  „Aggressive“ reduction also through lossy compression
- Tape storage budget has large uncertainties, but huge resources needed for sure!
  Ideas to reduce tape needs exist (eg. RAW compression, or not storing HITs), but impact still unclear

Storage Projections

Disk

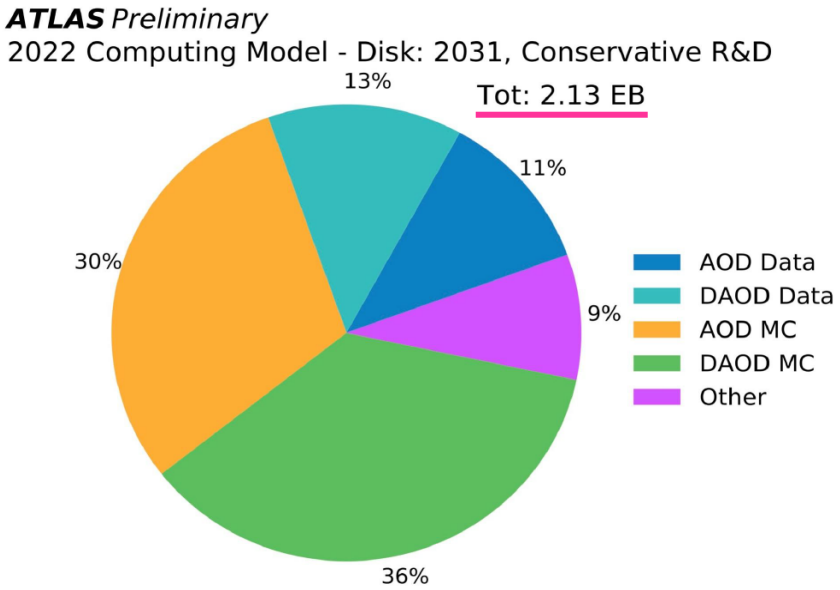

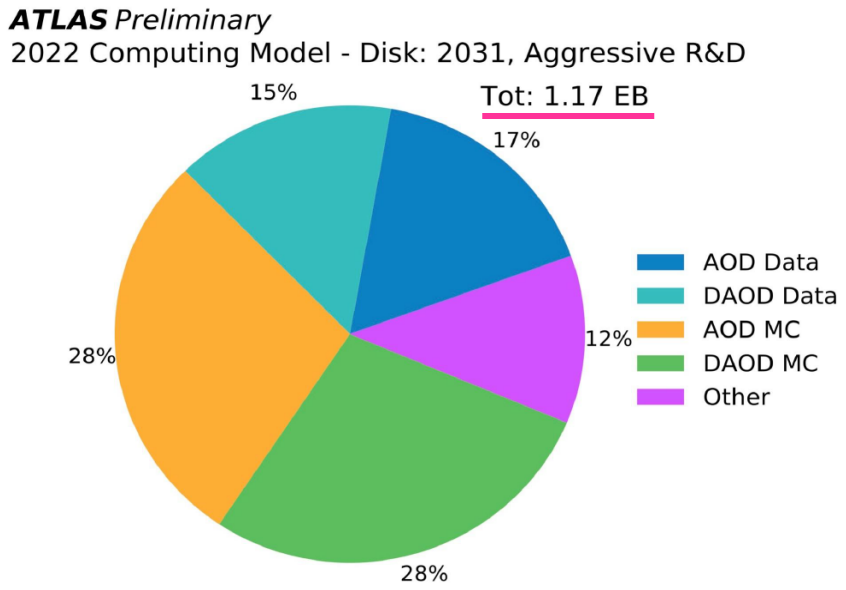

Tape

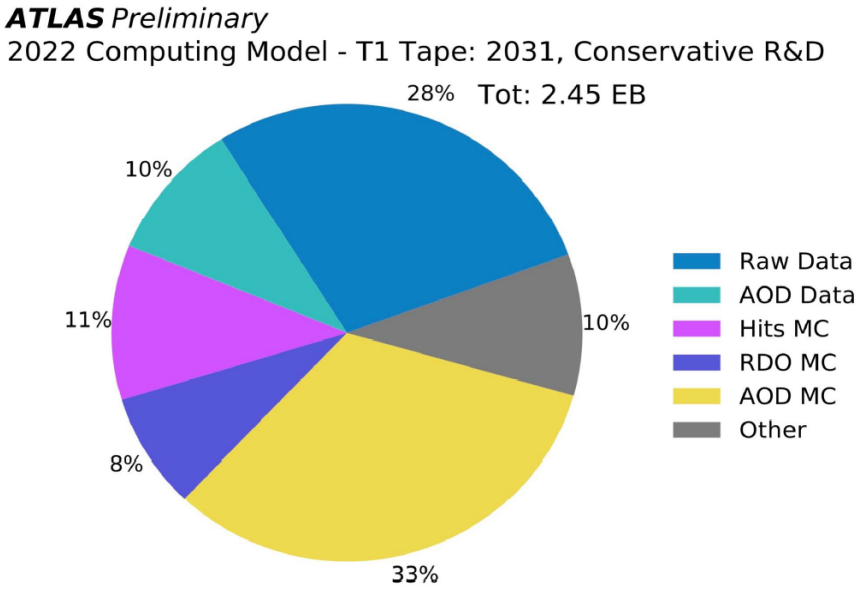

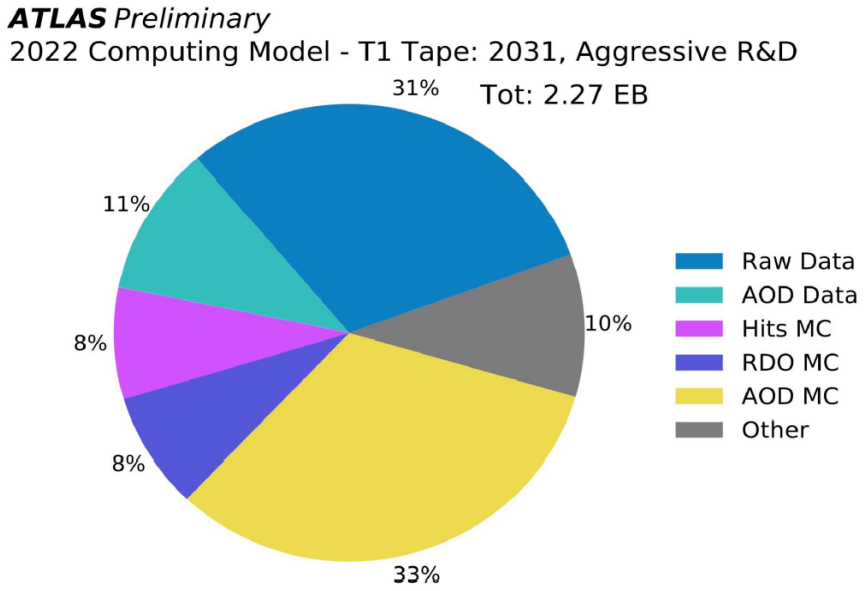

Fractions similar for both scenarios, but aggressive R&D can reduce disk storage by almost factor 2