

Interfaces

EIC Software: ToDo	#35
Responsible	T. Wenaus
Team	-
Review: Technical	-
Review: Documentation	-

Work

Comment on flat files

Comment on separating high-level logical data layout (C++ classes) from low-level physical data layout (columns of simple types)

Torre Wenaus - ATLAS experience broadly speaking has been 'flatter/simpler is better'. An outcome of the assessment of the analysis model in Run-1 was that the AOD analysis format should be simplified. In Run-1 it was usable only within the Athena framework; for Run-2 it was simplified and made usable directly in ROOT (as well as Athena), which proved a very popular move. More recently other considerations motivate simplification of both persistent and transient models. For both GPU and ML applications, I/O can dominate and overwhelm performance gains. The overheads are minimized if both persistent and transient representations map well to the simple structures favored by these applications. ATLAS is beginning to face this now. The transient event data model StoreGate in particular has never been reengineered, it was designed to be an object oriented whiteboard for physics data objects and is unsuited to injecting data onto a GPU with minimal transformation.