



Contribution ID: 558

Type: **Parallel Talk**

## Shared I/O Developments for Run 3 in the ATLAS Experiment

*Saturday, 9 July 2022 15:45 (15 minutes)*

The ATLAS experiment extensively uses multi-process (MP) parallelism to maximize data-throughput especially in I/O intensive workflows, such as the production of Derived Analysis Object Data (DAOD). In this mode, worker processes are spawned at the end of job initialization, thereby sharing memory allocated thus far. Each worker then loops over a unique set of events and produces its own output file, which in the original implementation needed to be merged at a subsequent step that would be executed serially. In Run 2, SharedWriter was introduced to perform this task on-the-fly, with an additional process merging data from the workers while the job was running, eliminating the need for the extra merging step. Although this approach had been very successful, there was room for improvements, most notably in the event-throughput scaling as a function of the number of workers. This was limited by the fact that the Run 2 version does all data compression within the SharedWriter process. For Run 3, a new version of SharedWriter has been written to address the limitations of the original implementation by moving compression of data to the worker processes. This development also paves the way for using it in a hybrid mode of multi-thread (MT) and MP workflows to maximize the I/O efficiency. In this talk, we will discuss the latest developments in Shared I/O in the ATLAS experiment.

### In-person participation

Yes

**Primary author:** METE, Alaettin Serhan (Argonne National Laboratory (US))**Presenter:** METE, Alaettin Serhan (Argonne National Laboratory (US))**Session Classification:** Computing and Data handling**Track Classification:** Computing and Data handling