

Metadata for fine-grained processing at ATLAS

Thursday, 13 October 2016 16:30 (15 minutes)

High energy physics experiments are implementing highly parallel solutions for event processing on resources that support concurrency at multiple levels. These range from the inherent large-scale parallelism of HPC resources to the multiprocessing and multithreading needed for effective use of multi-core and GPU-augmented nodes.

Such modes of processing, and the efficient opportunistic use of transiently-available resources, lead to finer-grained processing of event data. Previously metadata systems were tailored to jobs that were atomic and processed large, well-defined units of data.

The new environment requires a more fine-grained approach to metadata handling, especially with regard to bookkeeping. For opportunistic resources metadata propagation needs to work even if individual jobs are not finalized.

This contribution describes ATLAS solutions to this problem in the context of the multiprocessing framework currently in use for LHC Run 2, development underway for the ATLAS multithreaded framework (AthenaMT) and the ATLAS EventService.

Secondary Keyword (Optional)

Data model

Primary Keyword (Mandatory)

Data processing workflows and frameworks/pipelines

Tertiary Keyword (Optional)

Primary author: CRANSHAW, Jack (Argonne National Laboratory (US))

Presenter: CRANSHAW, Jack (Argonne National Laboratory (US))

Session Classification: Posters B / Break

Track Classification: Track 4: Data Handling