

ATLAS Distributed Computing experience and performance during the LHC Run-2

Tuesday, October 11, 2016 2:30 PM (15 minutes)

ATLAS Distributed Computing during LHC Run-1 was challenged by steadily increasing computing, storage and network requirements. In addition, the complexity of processing task workflows and their associated data management requirements led to a new paradigm in the ATLAS computing model for Run-2, accompanied by extensive evolution and redesign of the workflow and data management systems. The new systems were put into production at the end of 2014, and gained robustness and maturity during 2015 data taking. ProdSys2, the new request and task interface; JEDI, the dynamic job execution engine developed as an extension to PanDA; and Rucio, the new data management system, form the core of the Run-2 ATLAS distributed computing engine.

One of the big changes for Run-2 was the adoption of the Derivation Framework, which moves the chaotic CPU and data intensive part of the user analysis into the centrally organized train production, delivering derived AOD datasets to user groups for final analysis. The effectiveness of the new model was demonstrated through the delivery of analysis datasets to users just one week after data taking, by completing the calibration loop, Tier-0 processing and train production steps promptly. The great flexibility of the new system also makes it possible to execute part of the Tier-0 processing on the grid when Tier-0 resources experience a backlog during high data-taking periods.

The introduction of the data lifetime model, where each dataset is assigned a finite lifetime (with extensions possible for frequently accessed data), was made possible by Rucio. Thanks to this the storage crises experienced in Run-1 have not reappeared during Run-2. In addition, the distinction between Tier-1 and Tier-2 disk storage, now largely artificial given the quality of Tier-2 resources and their networking, has been removed through the introduction of dynamic ATLAS clouds that group the storage endpoint nucleus and its close-by execution satellite sites. All stable ATLAS sites are now able to store unique or primary copies of the datasets.

ATLAS Distributed Computing is further evolving to speed up request processing by introducing network awareness, using machine learning and optimization of the latencies during the execution of the full chain of tasks. The Event Service, a new workflow and job execution engine, is designed around check-pointing at the level of event processing to use opportunistic resources more efficiently.

Tertiary Keyword (Optional)

Secondary Keyword (Optional)

Primary Keyword (Mandatory)

Data processing workflows and frameworks/pipelines

Primary author: FILIPCIC, Andrej (Jozef Stefan Institute (SI))

Session Classification: Track 3: Distributed Computing

Track Classification: Track 3: Distributed Computing