

Sailing the petabyte sea: navigational infrastructure in the ATLAS event store

Monday, 13 February 2006 11:00 (20 minutes)

ATLAS has deployed an inter-object association infrastructure that allows the experiment to track at the object level what data have been written and where, and to assign both object-level and process-level labels to identify data objects for later retrieval. This infrastructure provides the foundation for opportunistic run-time navigation to upstream data, and in principle supports both dynamic determination of what data objects are reachable, and controlled-scope retrieval. This infrastructure is complementary to the coarser-grained bookkeeping and provenance management system used to identify which datasets were input to the production of which derived datasets, adding the capability to determine and locate the objects used to produce specific derived objects. It also simplifies the task of populating an event-level metadata system capable of returning references to events at any stage of processing. The tension between what the infrastructure can demonstrably support at site-level scales—it is already extensively utilized by ATLAS physicists—and what is expected to be constrained by policy in light of anticipated distributed storage resource limitations—is also discussed.

Primary authors: Dr SCHAFFER, Arthur (LAL ORSAY); Dr MALON, David (ARGONNE NATIONAL LABORATORY); Dr CRANSHAW, Jack (ARGONNE NATIONAL LABORATORY); Dr VAN GEMMEREN, Peter (ARGONNE NATIONAL LABORATORY)

Presenter: Dr MALON, David (ARGONNE NATIONAL LABORATORY)

Session Classification: Poster

Track Classification: Software Components and Libraries