

Organising scientific data by dataflow optimisation on the petascale

Tuesday, February 12, 2008 4:00 PM (0 minutes)

We analyse the Distributed Data Management system Don Quijote 2 (DQ2) of the High-Energy Physics experiment ATLAS at CERN. ATLAS presents unprecedented data transfer and data storage requirements on the petascale and DQ2 was built to fulfill these requirements. DQ2 is built upon the EGEE infrastructure, while seamlessly enabling interoperability with the American OSG and the Scandinavian NorduGrid infrastructures. Thus it serves as a relevant production-quality system to analyse aspects of dataflow behaviour in the petascale. Controlled data transfers are analysed using the central DQ2 bookkeeping service and an external monitoring dashboard, provided by ARDA. However monitoring dynamic data transfers of jobs and enduser data transfers cannot happen centrally because there is no single point of reference. Therefore we provide opportunistic clients tools for all scientists to access, query and modify data. Those tools report the needed usage information in a non-intrusive, scalable way.

3. Impact

We characterise three areas for improvement of dataflow. First, controlled data transfers issued by experiment operators or gridsite operators. This is constant data export from the experiment to distributed computing facilities, mostly defined by experiment computing models. Second, dynamic data transfers issued by jobs on a gridsite. Those production jobs may need to access data that is only available on remote sites. Third, uncontrolled data transfers issued by endusers; scientists fetching data for direct analysis. We argue that on the petascale complete replication of files is not a suitable option anymore as there is too much data and that erratic and unpredictable data movements are the norm. Furthermore it is important to value the relevance of certain data with respect to time to find useful data on the grid. Our model derives those usage patterns implicitly. Therefore global data movement and usage patterns on data must be taken into account when doing job/data co-allocation.

URL for further information:

<http://www.dps.uibk.ac.at/>

4. Conclusions / Future plans

The objective of reasonable organisation of scientific data on the grid is not a new one. Already, many approaches especially in file replication show good improvements. We argue though that once we approach petascale, low-level file reorganisation is not sufficient anymore and a global view of grid dataflow must be taken into account. We provide a preliminary model and its accompanying tools to understand erratic and unpredictable dataflows and show their usefulness in the production EGEE grid.

Provide a set of generic keywords that define your contribution (e.g. Data Management, Workflows, High Energy Physics)

Data Management, Dataflow, Grid Behaviour, Petascale

1. Short overview

Scientific applications on the grid are in most cases heavily data-dependent. Therefore, improving scheduling decisions based on the co-allocation of data and jobs becomes a primary issue. Hence, it is crucial to analyse the behaviour of existing data management systems in order to provide accurate information for decision-making middlewares in a scalable way. We show current research issues in understanding the behaviour of data management systems on the petascale to improve grid performance.

Primary author: Mr LASSNIG, Mario (CERN & University of Innsbruck, Austria)

Presenter: Mr LASSNIG, Mario (CERN & University of Innsbruck, Austria)

Session Classification: Posters

Track Classification: Existing or Prospective Grid Services